Version 7

Using Behavioral Big Data for Public Purposes: Exploring Frontier Issues of an Emerging Policy Arena

Rohan Samarajiva, Sriganesh Lokanathan

1st February 2016

The work was funded by New Venture Fund

Acknowledgements

The authors acknowledge the research assistance of Shazna Zuhyle, Dedunu Dhananjaya and in particular Chiranthi Rajapakse. The authors also express their gratitude for the input from researchers and subject experts who provided comments on the draft report through email, and through a workshop held in Colombo, Sri Lanka on 17th January 2016. These experts include Dr. Prabir Sen, Prof. Louiqa Raschid, Prof. Ryosuke Shibasaki, Dr. Ruhiya Seward, Prof Moinul Zaber, Prof. Ruvan Weerasinghe, Prof. Amal Kumarage, Dr. Amal Shehan Perera, and Helani Galpaya. In particular the authors are grateful for the engagement of Dr. Linnet Taylor who provided extensive and substantial comments on the draft report.

<Page left intentionally blank for use by the Foundation>

Executive Summary

Much of the discussion of the socio-economic implications of behavioral data has focused on the inclusion of more citizens and more aspects of their lives within the sphere of control enabled by pervasive data collection. Effective public policy rests on good information about problems and the efficacy of the deployed solutions. Governments obtained such information through National Statistical Organizations (NSOs) in the 19th and 20th Centuries. The modality in the 21st Century is the analysis of Big Data. Big Data will supplement old methods. It will also make it possible to have a fine-grained understanding that was not hitherto possible. The negligible incremental costs of analysis will make it possible to obtain insights more frequently, and even to conduct policy experiments.

This report explores privacy issues at the frontiers of research on and applications of behavioral big data for public purposes. It examines marginalization or exclusion from the scope of data collection. It also examines poverty/wealth mapping, including redlining, and the identification of regular and ad hoc congregations. In addition, it presents the state of the art on technical means used to mask PII in big data sets.

Marginalization has to be addressed as a special case of the problem of representivity. While representivity was neglected to some extent in the early days of big data, it is now receiving renewed attention. Marginalization is particularly important in the policy arena, where action could be triggered by insights that have been produced on the basis of available data. Even if unintentionally, people and problems not represented in the data could be neglected. Especially in the developing countries where datafication is incomplete, the question of representivity must be asked in relation to the questions being addressed.

There is much interest in correlating socio-economic data with geographic locations in the form of poverty mapping. Poverty mapping can help targeted delivery of services by government and other relevant agencies. But it is not possible to map poverty without also mapping wealth. Knowledge of where people of wealth are concentrated may result in those areas being prioritized for delivery of certain forms of services. This would, in most countries, be unlawful or politically unacceptable if done by governments or monopolistic suppliers acting under authority of government. Under conditions of competitive supply, some firms may choose to supply the wealthy areas while others may choose to concentrate on the poor areas. This would not be unlawful in most circumstances and may even be a feature of competitive supply.

Redlining, or the refusal to serve persons from specific geographical areas, is a phenomenon that has drawn the attention of policy makers in the United States. Historically, this has been done on the basis of crude correlations between location and ability to pay. In many cases this has also been correlated with ethnic identity. Big data may enable forms of discrimination more precise than those associated with redlining. Algorithms may be used to mask unlawful forms of discrimination. They may also lead to more accurate identification of consumers with desirable or undesirable characteristics or propensities and end the crude and error-ridden forms of discrimination known as redlining. Indeed, data analytics may enable first-degree price discrimination, displacing traditional ways of pricing products.

MNDB and other forms of big data that yield insights on movement of people through time and space can allow the identification of regular and ad hoc congregations in specific locations. Insights from pseudonymized historical data can be useful for deciding on locations of government and retail outlets and also for the pricing of outdoor advertising. Location-based advertising is of course a prime application. If location-based advertising is based on cell broadcasting the piercing of the collective shell is not necessary. However, other forms of location-based advertising will identify individual members of the congregation. When it comes to analysis of real-time and non-anonymized data problems emerge. Participants in political protests may be identified and acted against, posing serious issues.

Issues of collective privacy apply to both aspects discussed above. While some degree of harm may occur, it is concluded that it is not advisable to extend the concept of privacy, which is one that is applicable at the individual level, to the collective level. This will also negate most efforts to make efficient the delivery of public services.

In the case of technical methods of masking PII from individual data within big data sets, there is no easy solution though considerable advances have been made. In the case of developing countries, the current overall low levels of datafication offer some safeguards against re-identification of datasets with masked PII. Until more sophisticated technical solutions are found the data sets should be used in conjunction with non-technical safeguards such as legal agreements.

Table of Contents

List of Acronyms

BTS	Base Transceiver Station
CDR	Call Detail Record
FISA	Foreign Intelligence Surveillance Act
GR	Geographical Region
IGO	International Government Organization
MNBD	Mobile Network Big Data
MNO	Mobile Network Operator
NSA	National Security Agency
NSO	National Statistical Organizations
PII	Personally Identifiable Information
SEL	Socio Economic Levels
SIM	Subscriber Identity Module
VLR	Visitor Location Registry

Introduction

This report explores privacy issues at the frontiers of research on and applications of behavioral big data for public purposes. The focus of the present discussion is on the subset of big data known as transaction-generated data (also described as "data exhaust") arising from the day-to-day behaviors of persons and the technological devices closely associated with them.

Effective public policy rests on good information about problems and the efficacy of the deployed solutions. Governments obtained such information through National Statistical Organizations (NSOs) in the 19th and 20th Centuries. The modality that will increasingly be used in the 21st Century is analysis of Big Data.

Big Data will supplement old methods. It will also make it possible to have a fine-grained understanding that was not hitherto possible. The negligible incremental costs of analysis will make it possible to obtain insights more frequently, and even to conduct policy experiments. As citizens engage in various activities such as interacting with large organizations, making phone calls, consuming electricity, and even just moving around, they generate large volumes of datafied[1] records, which may be analyzable depending on the extent of computerization of the systems they interact with. Advances in computer memory and software have made it easier to analyze these vast volumes of data and extract policy-relevant insights. Modern governments seek to exploit this potential for public purposes.

Since 2012, LIRNE*asia* has been engaged in research on Mobile Network Big Data (MNBD). At this time, MNBD is the only big data set that has representivity adequate for public policy problems in the areas of urban development, transport planning, socio-economic monitoring and epidemiology in the developing world (Samarajiva, et al., 2015).

In parallel with this research LIRNE*asia* has been studying the associated issues of implications for competition, privacy and marginalization, with the objective of reducing the transaction costs of releasing data to third-party researchers by Mobile Network Operators (MNO) who collect and control the data. While the frontier issues addressed in this report are not essential for the immediate purpose of reducing the transaction costs of releasing data to third-party, public-interest researchers, they could become centrally relevant as the field develops. For example, understanding the emerging technical solutions to the problem of effectively masking personally identifiable information (PII) is important for the development of "future-resilient" elements that could be used in non-disclosure agreements, guidelines and regulations.

Big data and the problem of control

Information and control are closely connected. Beniger (1986, pp. 7-8) states that the twin activities of information processing and reciprocal communication (or feedback) are inseparable from the concept of control. Control is defined in the broadest sense as "purposive influence toward a predetermined goal." Even though he wrote well before the current democratization of big data analytics, Beniger provides possibly the best answer to the questions "why big data?" and "why now?"

[1] "Datafication" is defined as transforming a phenomenon into a quantified format that allows it to be measured and analyzed: Mayer-Schonberger & Cukier (2013): pp. 78-86.

Public policy is necessarily intertwined with issues of control that can range from "hard' or "soft" control of behavior to the control of undesirable forms of control of one group in society by another. Much of present-day concerns about the negative effects on privacy are based on the perceived increase in the gathering of data that could lead to greater control; about more aspects of citizen's lives being made visible to governments or to corporations. This focus on privacy problems associated with inclusion within the sphere of control will be dominant in this report, though non-inclusion or marginalization is also addressed. Given the centrality of privacy, the analysis of the frontier issues will be preceded by an explication of privacy in a manner conducive to translation into policy-relevant form.

The four frontier areas that are discussed are marginalization, the implications of poverty mapping, including redlining, and of identifying congregations and technical methods of masking identity.

The insights gained can be used to develop elements that may be included in guidelines, codes of conduct or legal agreements governing large data sets that adequately represent populations relevant for public-policy purposes. They can also help broaden and illuminate the discourse on the social implications of big data.

Privacy

Privacy, as commonly understood, "is a sweeping concept, encompassing (among other things) freedom of thought, control over one's body, solitude in one's home, control over personal information, freedom from surveillance, protection of one's reputation, and protection from searches and interrogations" (Solove, 2008, p. 1). Attempts to define it in terms of boundary control by individuals (e.g., Samarajiva, 1994: 90) are difficult to translate into practical policy. For example, it is difficult to clearly demarcate what an individual has authority over in the case of data generated as a by-product of a transaction, where the data are co-produced and held by one party.

Solove (2008, p. 174) argues that privacy as an abstract concept is difficult to pin down, because it "involves a cluster of protections against a group of different but related problems." He concludes, correctly, that the focus should be shifted away from defining privacy, to addressing privacy problems (or harms). He proposes 16 privacy problems, grouped into four general types: Information collection (comprising surveillance and interrogation); information processing (comprising aggregation, identification, insecurity, secondary use and exclusion); information dissemination (comprising breach of confidentiality, disclosure, exposure, increased accessibility, blackmail, appropriation and distortion); and invasion (intrusion and decisional interference) (Solove, 2008, Ch. 5). Harms that may be caused by behavioral big data or transaction-generated data that fall within the scope are primarily located in the second of the clusters, information processing, and secondarily in information collection, the first cluster, and information dissemination, the third cluster.

Surveillance

Within the information-collection cluster proposed by Solove, the most relevant problem is surveillance. In the context of behavioral big data, it is useful to distinguish between active and passive surveillance. Installation of a device such as a GPS tracker constitutes active surveillance.[2] Active surveillance, where the activity is undertaken for the primary purpose of collecting data on a specific individual is normally associated with law enforcement and espionage and is, for the most part, a "small data" problem. What is relevant in the context of big data is passive surveillance in the form of data that are a by-product of some activity (Mundie, 2014). Where systems are explicitly engineered to collect more data than are needed for normal operations, the line between passive and active is blurred.[3]

The harms are the gathering of information about a person through active or passive surveillance. The former may be prohibited or constrained. But the latter is difficult to control without stifling the activity that generates the data as by-product. If the base activity is one that benefits the data subject and is one that he/she engages in willingly, there may be merit in not prohibiting collection, and instead focusing remediation on subsequent processing, as suggested by Mundie (2014).

---

2 *United States v. Jones*, 132 S. Ct. 945, 565 U.S. (2012).

3 The US Communications Assistance to Law Enforcement Act (CALEA) of 1994 is one of the earliest examples involving electronic technology. http://itlaw.wikia.com/wiki/CALEA

Aggregation[4]

Aggregation, as defined by Solove (2008), can take two principal forms in relation to behavioral big data. First, it is the aggregation of discrete data elements related to a single individual within one dataset, e.g., not just the datum that A interacted with B, but the pattern of A's interactions with B and vice versa. Second is the aggregation of data from different sources, e.g., from mobile networks and from surveys or from payment terminals in shops. Pseudonymization is not a barrier to the aggregation of data regarding a person within a dataset, though the resulting insights about the digital person will not be connected to the person in "realspace." Pseudonymization makes aggregation across multiple data sets more difficult.

Aggregated data yields a richer picture than non-aggregated data. Aggregation may also reduce the potential for wrong conclusions being drawn from the partial picture presented by non-aggregated data.[5]

Therefore, the first set of potential harms comprises errors caused by aggregation or lack thereof. The second is about "true" insights drawn through aggregation, when the "truth" is not intended to be disclosed. The third is about the dangers of identification through de-anonymization made possible because of aggregation. At the individual level, the third is the most significant.

One may ask what harm is caused by erroneous or "truthful" information generated through aggregation as long as the data subject is anonymous. So for example, one may conclude through aggregation that a particular data subject has undergone an illegal/morally questionable medical procedure. This may be true, or may be false because the aggregation was incomplete and missed some significant data (the data subject may be visiting the medical facility for a different reason). As long as the data subject cannot be identified, it is difficult to discern the harm at the individual level.

However, harm may occur to an organization or a group using that organization's services. It may be possible to infer the location of an illegal service provider using aggregated anonymized/pseudonymized data sets even if the identities of individuals using the services continue to be effectively masked. While the specific persons included in the data sets may escape prosecution, the organization providing the service and future users may suffer the consequences of engaging in actions illegal under that country's laws. Increasingly, law enforcement authorities are using analytics for purposes such as predictive policing.[6] Whether we describe the consequences of such actions as harmful or not depends on the purpose. If against criminals or those engaging in socially undesirable actions, it is unlikely that it will fall within the definition of harm as discussed here.

---

[4] The term aggregation is here used not as a tool for obscuring identity as it is sometimes understood, but exactly in the sense used by Solove (2008). It is a technical term that is central to his analysis.

[5] Recognizing, of course, that all data are partial representations of "reality." The debate is not about fully accurate versus inaccurate, but about the relative veracity of partial representations.

[6] Perry, W.L.; McInnis, B.; Price, C.C.; Smith, S.C.; Hollywood, J.S. (2013). *Predictive policing*. Santa Monica CA: Rand.

At the individual level, the harm is in the likelihood that aggregation may permit identification through de-anonymization. At the group level, it is possible that harm may result if techniques used in law enforcement are used against political actors.

Identification, individual and group

Identification is a central concept. According to Solove (2008: 122-25), identification "is connecting information to individuals. . . . Aggregation creates . . . a portrait composed of combined information fragments. Identification goes a step further—it links the digital person directly to a person in realspace."

It is clear that identification is an essential element of the postulated harms at the individual level, where much, if not all, of the privacy discussions focus. But it is also the essential element in harms at the collective or group level (discussed below).

Insecurity

"Glitches, security lapses, abuses and illicit uses of personal information all fall into this category [of] insecurity, . . . a problem caused by the way our information is handled and protected" (Solove, 2008, p. 127). As the volume and value of aggregated data increases (becoming big data), the harms that can be caused by the data falling into wrong hands or being distorted increase. Here too, the harm at the individual level is tied to identity. Effectively anonymized data falling into the hands of an ill-meaning or unintended person or organization is unlikely to cause a person whose data are included within the data set any harm.

However, some scholars such as Taylor (2015) contend harms may be caused to groups from anonymized data falling into the hands of unintended persons.

Secondary use

"'Secondary use' is the use of data for purposes unrelated to the purposes for which the data was initially collected without the data subject's consent" (Solove, 2008: p. 131). The definition hints that it is an artifact of law developed in the 1970s anchored in practices such as individuals filling out forms and ticking boxes indicating consent that have little relation to the passive and pervasive surveillance that is the norm today. When one makes a phone call, one generates a Call Detail Record (CDR). Was the data given or collected, or was it jointly generated in the course of completing the call? How and when could consent be given? Is it possible to maintain an effective mobile network without aggregating and analyzing different elements of data within the CDR such as the loading of the Base Transceiver Station (BTS)? Is the use of the data for network optimization a secondary use?

Secondary-use absolutism poses the danger that uses by all but the entity co-generating the data will be prohibited. As senior Microsoft official Craig Mundie (2014) states "today, there is simply so much data being collected, in so many ways, that it is practically impossible to give people a meaningful way to keep track of all the information about them that exists out there, much less to consent to its collection in the first place."

One way this problem may be managed is through omnibus consent forms that may be obtained at the moment of establishing the commercial relationship. Depending on the skill of the lawyers drafting the documents, one would have to give consent to all imaginable uses by the provider of goods or services, or make do without the service.[7] Since this particular subterfuge will not be effective in the case of third parties, the practical result will be exclusion of all third parties from the benefits of data analytics of data co-generated by others. In the case of for-profit entities, the loss will be to innovation and competition. The use of big data for public purposes will also suffer.

Exclusion

Solove (2008, pp. 134-35) proposes the term "exclusion"[8] for failure to provide individuals with notice and input about their records. He states that the harm is created by the data subject being shut out from participating in the use of the data, from not being informed about how it is used, and by not being able to affect how it is used. While it is present in Fair Information Practices, Solove (2008, p. 207) states that "for the most part, tort law has not recognized exclusion as a harm,"

The Kafka quotation used by Solove (2008, p. 133) illustrates the possible harm: "For in general the proceedings are kept secret not only from the public but from the accused as well." When benefits/harms are decided on the basis of data sets, the argument is that not only the data but the algorithms that are used to extract insights from them must be known and subject to correction (Pasquale, 2015; Tufekci, 2014).

Concern about exclusion or opacity is intuitively correct for credit reports, the starting point of modern privacy remedies. But the harms are small compared to the massive transaction costs that would be associated with notifying all data subjects whose data are in big data sets and permitting them rights to examine and correct them. For example, every BTS in a mobile network contains data on thousands of "data subjects" including ephemeral data as such as what is recorded on the Visitor Location Registry (VLR) on when they moved within the range of the BTS and when they moved out. It would serve little purpose to notify them of this. The transaction costs would be very high. Allowing access to commercially sensitive data sets would also not be practical.

The algorithms applied to the data to produce insights pose difficulties of a higher order of magnitude. Even if the data were understandable, there are few realistic solutions to the problem of eliminating the opacity of the algorithms (Pasquale, 2015, ch. 6).

Exclusion, therefore, poses no harm in relation to many forms of transaction-generated big data such as MNBD. It could, however, be the cause of considerable problems in the form of high transaction costs if attempts were made to apply remedies that may have been appropriate in the days of credit reports.

---

[7] "Because privacy notices under the 1980 Guidelines constrain future data uses, notices have become increasingly broad and permissive. The result has been the increasing erosion of information privacy." –Cate, Cullen & Mayer-Schonberger (2013).

[8] Perhaps the least felicitous of the set.

Breach of confidentiality

Most privacy problems sought to be addressed by the tort of breach of confidentiality are not relevant to big data such as MNBD. It requires consideration because of the "third-party doctrine" exemplified by the *United States v. Miller* and *Smith v. Maryland* decisions which govern government access to transaction-generated data of individuals (small data).[9] In the former, the US Supreme Court held that no breach occurred when a person's bank records were released to government because "all of the documents obtained, including financial statements and deposit slips, contain only information voluntarily conveyed to the banks and exposed to their employees in the ordinary course of business."[10] In *Smith v Maryland*, the logic was extended to call details (not the content of the call), on the basis that people "know that they must convey numerical information to the phone company," and, cannot "harbor any general expectation that the numbers they dial will remain secret."[11]

The US government's justification for the collection and use of telephone metadata pertaining to US citizens by the National Security Agency (NSA) exposed by Snowden was based on the third-party doctrine, derived from the above judgments (Savage, 2013). A 2013 decision from the District Court of the District of Columbia (perhaps the most important, because Washington DC is within the District) attracted significant attention because it explicitly contradicted the Smith rationale, stating that the surveillance of meta data in 2013 was qualitatively different from that which was decided in 1979.[12] However, a subsequent decision by a District Judge from the Foreign Intelligence Surveillance Act (FISA) Court responsible for oversight of the National Security Agency's surveillance activities reaffirmed the third-party doctrine. Until the various appeals work their way up to the Supreme Court, *Smith v Maryland* will continue as the ruling precedent in the US. As stated by the FISA judge: "The Supreme Court may someday revisit the third-party disclosure principle in the context of 21st-century communications technology, but that day has not arrived" (Savage, 2013)."

It must be noted that there is no question in either *Miller* or in *Smith* about whether the bank and the telephone company could use the data. The only question at issue was whether the data could be given to a third party, the government, without the data subject's authorization. Since the focus here is on use of transaction-generated data by third parties, the privacy problem or harm may be restated as one of harms cause by aggregation and identification at the individual or collective levels, as discussed above.

Disclosure

Disclosure refers to disclosure of true information about a person. In some countries, there are laws restricting the disclosure of data from educational institutions, video rental companies, health services, etc. The harm caused by disclosure is damage to reputation. Reputation being

---

9 425 U.S.435 (1976) and 442 U.S. 735, respectively.

10 425 U.S. 435 (1976), at 442-43.

11 442 U.S. 735 (1979), at 743.

12 Klayman v Obama, Civil Action 13-0851(RJL).
http://www.nytimes.com/interactive/2013/12/17/us/politics/17nsa-ruling.html?ref=politics&_r=0

tied to identity, anonymization can avoid the harm at the individual level. There may be circumstances under which groups suffer harm, but they have to be dealt with on a case-by-case basis, outside the realm of privacy.

Increased accessibility

Here, the information is public, but is difficult to get to. This is an important issue in the context of the Internet, with its easy search capabilities, and the increasing trend toward open data and open government. It primarily applies to public records held by government and not to data held by private entities where there is no presumption of openness.

But the issue may become relevant if and when data such as MNBD in raw or semi-processed form are made available on the web, especially if these actions are a result of government direction.[13]

[13] For a discussion in the context of open government, see Borgesius, van Eechoud, & Gray (2015).

Marginalization

Lerman (2013) sketches out two archetypes relevant to big data analytics, and extends to a third:

> The first is a thirty-year-old white-collar resident of Manhattan. She participates in modern life in all the ways typical of her demographic: smartphone, Google, Gmail, Netflix, Spotify, Amazon. She uses Facebook, with its default privacy settings, to keep in touch with friends. She dates through the website OkCupid. She travels frequently, tweeting and posting geotagged photos to Flickr and Instagram. Her wallet holds a debit card, credit cards, and a MetroCard for the subway and bus system. On her keychain are plastic barcoded cards for the "customer rewards" programs of her grocery and drugstore. In her car, a GPS sits on the dash, and an E-ZPass transponder (for bridge, tunnel, and highway tolls) hangs from the windshield.
>
> . . . . . .
>
> Now consider a second person. He lives two hours southwest of Manhattan, in Camden, New Jersey, America's poorest city. He is underemployed, working part-time at a restaurant, paid under the table in cash. He has no cell phone, no computer, no cable. He rarely travels and has no passport, car, or GPS. He uses the Internet, but only at the local library on public terminals. When he rides the bus, he pays the fare in cash.
>
> Today, many of big data's tools are calibrated for our Manhattanite and people like her—those who routinely generate large amounts of electronically harvestable information. A world shaped by big data will take into account her habits and preferences; it will look like her world. But big data currently overlooks our Camden subject almost entirely. (And even he, simply by living in a U.S. city, has a much larger data footprint than someone in Eritrea [the third archetype], for example.)

Lerman's short piece on exclusion is an exception to the general emphasis on problems of inclusion. In many fields of public policy, practitioners are well aware of the problem of exclusion, such as that of those who administer sample surveys oversampling roadside communities and excluding those who are more difficult to reach, and post-disaster aid not reaching those in less visible locations.

Box 1: Boston's Street Bump app

The City of Boston makes available an app called Street Bump that can be downloaded to smartphones. Any citizen can place the smartphone in a holder in a car and press one button to start the app at the beginning of a journey. No calls would be taken during the journey. The accelerometer of the smartphone collects data that has been proven to be effective in identifying pot holes and speed bumps. At the end, another button is pressed and the collected data including the GPS coordinates of the starting and ending points are sent to City Hall. Using algorithms the bumps that should be there and those that should not be there are identified and the latter get routed into the work order system for repairs.[14]

The assumption is that smartphones are ubiquitous in Boston. What if a similar crowdsourced big-data application is deployed in a city which has less than 10 percent smartphone users?

[14] http://www.cityofboston.gov/DoIT/apps/streetbump.asp

The issue has to be situated within the larger problem of representivity (Miller, et al., 2015; Samarajiva, 2014). Miller, et al. propose an approach that would require researchers to explicitly address the representivity of a particular data set in the hope that over-broad claims will not be made for it and biased policy prescriptions that would not be derived from the findings. Samarajiva, et al. (2015) argue for reliance on the less rich data generated by mobile networks (as against smartphones) in developing countries to avoid marginalizing the poor. The outcomes of marginalization may be optimal in terms of privacy because none of the privacy harms are caused by marginalization. Indeed, marginalization may well describe the aspiration of the privacy absolutists.

It must be noted that marginalization is not a binary condition, but that there is a continuum of conditions. Certain groups such as the homeless or illegal immigrants are marginalized by conventional surveys and censuses. MNBD cover more people than data collected from smartphones or from Twitter, but do not cover every person.

Implications of poverty and wealth mapping

Socio-economic mapping identifies the poor so services may be efficiently delivered to them. Thus, it is desired by some governments and international government organizations (IGO). At the present time, socio-economic mapping seeks to literally map or associate poverty on spatial representations. In the future, it may be extended beyond mapping in the literal sense. The analogy is to the zip-code-based voter mobilization efforts of past US elections versus the precision-targeted get-out-the-vote exercise of the 2012 Obama campaign. The discussion of collective privacy below addresses some of those broader concerns.

Frias-Martinez, et al. (2012) developed a mathematical model that computes approximate Socio Economic Levels (SEL) based on human-mobility variables derived from mobile network big data. Household survey data was used to determine SELs for each Geographical Region (GR) defined by the National Statistical Organization (NSO). In order to study the relationship between SELs and human mobility, they first geographically mapped the BTS coverage areas within the GRs, and computed a SEL value for each area of coverage of each BTS. Their results indicated that populations with higher SELs are strongly linked to larger mobility ranges than populations with lower socio-economic status. By extending this method, it is possible to create a model to estimate SELs based on mobile network big data.

Another study by Gutierrez, Krings, & Blondel (2013) used two types of mobile network big data, namely the communication network of the subscribers and history of airtime credit purchases to assess the SELs. The authors hypothesized that user who make large airtime purchases would be richer than those who make multiple small ones, as the poorer will not have enough ready cash to buy lots of airtime credit all at once. They combined this analysis with the study of social networks of the users (two users are considered as connected if they communicated with each other at least once during the month). The authors found that people with similar purchasing power tend to be connected.

Operators have better measures that could allow for poverty mapping as well as modeling economic shocks (Figure 1). Mobile network operators have real time measures of revenue at the BTS level. Emergent research in Africa is showing how these could also be used to model economic shocks (David, 2013). However given the sensitivity of revenue data for operators, such data are not generally available to outside parties.

Using MNBD (specifically the communication patterns as well as the history of airtime credit purchases) from the MNO, Orange, in Côte d'Ivoire, researchers estimated the relative income of individuals, and the diversity and inequality of income. They then used these measures to understand socio-economic segregation at a fine-grained level for Côte d'Ivoire. The above figure shows the poor areas (in blue) in relation to the areas of high economic activity (yellow to red areas). Giving due credit to the researchers who are breaking new ground with studies such as this, care must always be taken to validate the explanations with knowledge from multidisciplinary teams cognizant of ground conditions.

Figure 1: Poverty mapping in Côte d'Ivoire[15]

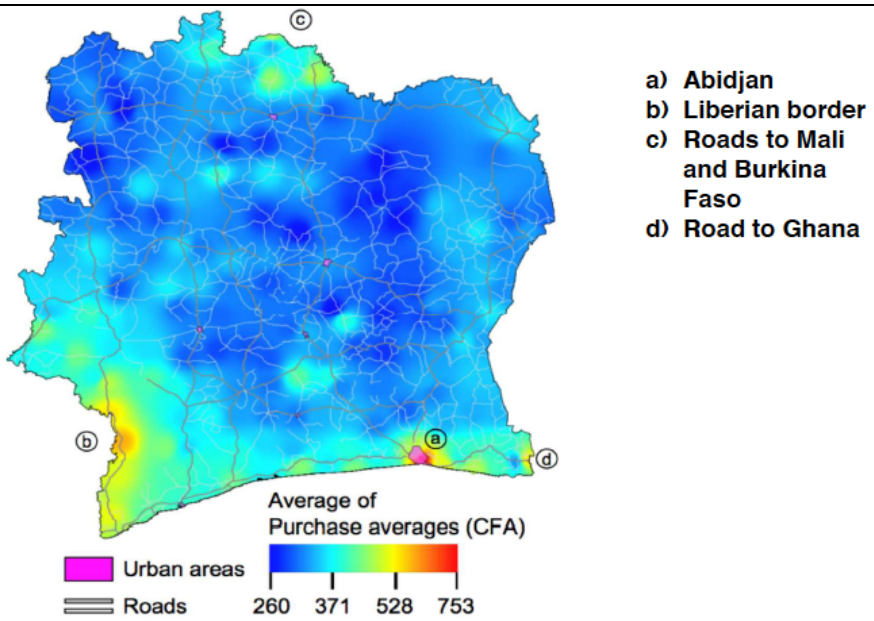

Using mobile network data (specifically the communication patterns as well as the history of airtime credit purchases) from Orange in Côte d'Ivoire, researchers estimated the relative income of individuals, and the diversity and inequality of income. They then collectively used these measures to understand socio-economic segregation at a fine-grained level for Côte d'Ivoire. The above figure shows the poor areas (in blue) in relation to the areas of high economic activity (yellow to red areas).

If the poor can be identified, it follows that the rich can too (as illustrated by Figure 1). Will this result in prioritization of the areas where the rich live in terms of service delivery, for example in terms of rolling out 4G networks or locating bank automatic teller machines?

In competitive markets, suppliers are not expected to serve the entire market at the very outset or even at any point. Uncertainty about demand is normal. Therefore, suppliers enter in limited geographical areas or focus on particular market segments at the outset. It is only on the basis of feedback from these activities that the firm will scale up. Some firms will adopt niche strategies and never seek to serve the entire market.

However, expectations are different for governments and private monopolies operating under license from government. Here, the supplier of services is obliged to provide service to all on a non- discriminatory basis (or, at least, strive to do so). Here, poverty or other kinds of mapping could be used either by the suppliers to reach desired groups or by regulators to ensure that they have not engaged in undue discrimination.

In the United States, perhaps because of the legacy of difficult race relations, discrimination known as "redlining" has been found to exist even in ostensibly competitive markets (Podesta et al., 2014, p. 53). Well before big data, some companies were using rough-and-ready indicators as well as rudimentary forms of data analytics to discriminate against or redline communities,

[15] Source: Gutierrez et al. (2013)

usually those of ethnic minorities. Government and civil society were also using data in their efforts to prevent forms of discrimination that are against the law.

It may be said that denying a person credit on the basis of where she lived (an illegal act in the US) is based on an algorithm, albeit a very rudimentary one that draws a direct correlation between a geographical area associated with a zip code and the likelihood of repaying the loan. Big data analytics offers the possibility of using more sophisticated algorithms.

Two possible trajectories exist. The first is that effort will be put into unlawful discrimination in ways that are difficult, if not impossible, to regulate because of the opacity of the deployed algorithms. The second is that the improved algorithms will enable accurate decision-making including discrimination (not all forms of discrimination are unlawful) without the mistake-laden, clumsy practices of the past. The first assumes that political and cultural imperatives will override economic incentives; the second sees the economic incentives as preeminent.

Algorithms are necessary for both trajectories. They will be opaque, as are all algorithms to those who do not have specialized knowledge. Algorithms in the first case would be made purposely difficult to understand, because disguising the illegal discrimination is one of its design parameters. As the recent Volkswagen scandal illustrates, software can be designed to deceive (Ewing, 2015). Those in the second case would not have been designed with that objective.

Coarse correlations between entitlements to credit or whatever else do not require the collective form to be pierced and direct relationships established between the digital person and the person in "realspace." A person can be denied credit based on the zip code of the area they live in, without even knowing the person's name. But the possible future trajectories, both negative and positive, require that such relationships be established because the correlations will no longer coarse.

One new issue related to redlining has been highlighted by Podesta (2014, pp. 46-47). The practice of some offline merchants offering different prices to those who fit different socio-economic profiles or come from different locales based on algorithms has been documented. The practice is already prevalent amongst online merchants. Valentino-Devries (2012) showed how online merchants displayed different prices to customers browsing their website based on the site's estimate of the customer's geographic location. The exact formula used to set prices was not clear; however the strongest correlation seemed to be the distance to a rival store. Hannak, et al. (2014) found multiple factors affecting the price offered.

In defense of this practice, it has been stated that the algorithms set prices based on availability of alternative suppliers. The offering of low prices to those living in rich neighborhoods is sought to be explained in terms of more alternatives being available to them. However explainable this may be from a theoretical perspective, it may be difficult to defend in the context of the values that govern public policy. Valentino-Devries (2012) reports that consumers consider such a practice as 'unfair.'

At a more abstract level, the problem is one of first-degree price discrimination. First-degree price discrimination, or person-specific pricing, has not been practiced or observed because it was not possible to discern reservation values. This constraint may be in the process of being

overcome now that capabilities exist to analyze individual behavior as recorded in multiple transaction-generated data sets (Shiller, 2014). Big data and electronic commerce have reduced the costs of targeting and first-degree price discrimination. It is argued that the increased availability of behavioral data may encourage a shift from third-degree price discrimination towards personalized pricing (Executive Office of the President of the United States, 2015).

Shiller opines that first-degree price discrimination that used to be taught as an abstract, but unrealizable, pricing strategy may soon become commonplace. While it may raise profits when implemented by firms in oligopolistic and differentiated-product markets, he posits that such outcomes may not occur when multiple firms implement it in competitive markets. Referring to the finding by Kahneman, et al. (1986) that an overwhelming majority of the public (91 percent of respondents in that study) perceived it to be unfair, Shiller foresees problems in public acceptance and thereby in how it is treated in public policy.

Crude forms of poverty or wealth mapping do not lead to privacy harms at an individual level. But the more sophisticated forms that are emerging will, because they require the piercing of the collective shell.

Implications of identifying congregations

Certain datasets such as those generated from mobile networks, smartphones and other mobile devices and stored-value cards used for public transport permit the generation of fine-grained insights about temporal movements of people. Using historical, pseudonymized MNBD it is possible to identify not only where people congregate, but also at what times in general. Using real-time analysis of MNBD in the form of CDRs generated for billing purposes or VLR data generated when terminal devices notify their presence within the signal area of a BTS, it would also be possible to identify irregular or ad hoc congregations as they occur.

Congregation patterns derived from pseudonymized data sets that include data on movement through time and space have many public and private applications. Knowing where people congregate is useful for deciding on where to locate government citizen-service centers as well as retail outlets and customer service centers. The New York City Office of Data Analytics is currently offering reports that draw from both private and public data sets to assist small businesses make informed decisions on where to locate: "detailed information about economic activity, demographics, foot traffic and other key business metrics around locations they are considering."[16] Other potential uses are rearranging pedestrian and other traffic patterns at different times of day (including experimentation), time-based pricing of outdoor advertising displays, etc.

The above applications are based on regular patterns and do not necessarily permit the persons whose aggregate movements are captured in the data to be directly identified or communicated with. However, if the data are not pseudonymized, it will be possible to identify those who are regularly in specific locales, for example for the targeted dissemination of location-based advertising.

Real-time analysis poses greater problems. Generally, real-time analysis will have to be done without masking the identities of the persons congregating, because pseudonymization is an additional procedure that would take time. The prime application being location-based advertising, preventing the advertiser from reaching the prospect would be counter-productive in any case.[17] Depending on whether the recipients of the location-based advertising perceives the messages as helpful or annoying, privacy complaints are likely to arise. Unless the concerns are handled carefully, productive uses of the capability to identify congregations for public purposes as well as for location-based advertising may be stymied.

Techniques used to enable location-based services can also be used to track movements of groups or individuals for other purposes. Indications exist that mobile network big data are already being used by governments to identify and control gatherings. In January 2014 text messages warned protestors in Kiev, Ukraine, that they were participants in a mass riot: "Dear

---

[16] http://www.nyc.gov/html/analytics/html/initiatives/economic_development.shtml

[17] Cell broadcasting, wherein all mobile devices within the coverage area of a BTS receive messages in broadcast mode without the addresses/numbers having to be known, allows for a form of location-based advertising without the conclusions drawn at the collective level having to be translated to the individual level. However, cell broadcasting lacks the precision of targeting that would be possible if the address/number were known and does not permit aggregation of transaction data over time.

subscriber, you are registered as a participant in a mass riot." The mobile operators MTS and Kyivsta issued statements claiming they were not responsible for the messages. The language of the texts were reported to echo the wording of new laws on public gatherings (Walker, 2014; Lopez, 2014). The capture of the mobile numbers is believed to have happened through a fake base station placed by government network hackers. It is believed that operators had refused to provide access to their networks.[18]

Social media such as Twitter can be powerful tools in organizing gatherings and protests particularly in situations where governments have censored mainstream media. The widespread use of Twitter was seen during demonstrations in Turkey in 2013 (Arsu, 2014, Parkinson, 2013). Thus, it would be natural to turn to social media data to predict (and prevent) gatherings, in instances when they were not shut down (Burns, 2011).

Kallus (2014) describes how big data on social media can be used to predict events. The author attempts to predict the occurrence, specific timeframe, and location of actions before they occur based on public data collected from over 300,000 open content web sources. The sources ranged from mainstream news to blogs and social media. Natural language processing was used to extract event information from the content. Statements made on Twitter about a future date from the time of posting were found to be particularly indicative. Botta, et al. (2015) conclude that "accurate estimates of the number of people in a given location at a given time can be extrapolated from mobile phone or Twitter data."

Box 2: Cutting edge of known research

Security agencies appear to be increasingly turning to social media surveillance to predict gatherings. One such initiative is the Open Source Indicators project run by the Intelligence Advanced Research Projects Activity (IARPA), USA, which aims to "develop methods for continuous, automated analysis of publicly available data in order to anticipate and/or detect significant societal events, such as political crises, humanitarian crises, mass violence, riots etc."[19]

The Embers (Early Model Based Event Recognition using Surrogates) project developed by Virginia Tech uses "open-source indicators"—social media, satellite imagery and more than 200,000 blogs. The project seeks to identify patterns that predict events such as civil uprisings, humanitarian crises, mass migrations, protests and riots. A Newsweek (2015) report states that the project was first used to examine open-source data streams in Latin America in relation to the World Cup protests in Brazil in 2013, and the violent student protests in Venezuela in 2014. It is now moving beyond Latin America to the Middle East and North Africa, covering countries such as Iraq, Syria, Egypt, Bahrain, Jordan, Saudi Arabia and Libya (Goodman, 2015)

Other projects run by the IARPA seek to match online and offline "behavioral indicators," including "ideology or worldview," and to extracts geolocation information from posts, photos, and videos. (Gould-Wartofsky, 2015).

Identifying congregations could lead to privacy harms, albeit at a collective level in the first instance. The harms and what may be done about them are discussed under collective privacy

[18] Personal communication, Linnet Taylor.

[19] http://www.iarpa.gov

below. However, when the digital persons and persons in realspace are linked, privacy harms become relevant.

Collective privacy

Integrally connected to the above discussion is the notion of collective privacy. Poverty and wealth mapping and the identification and attribution of characteristics to congregations involve collectives of different forms.

Identification is central to all discussions of privacy. Identification "is connecting information to individuals. . . . Aggregation creates . . . a portrait composed of combined information fragments. Identification goes a step further—it links the digital person directly to a person in realspace" (Solove, 2008, pp. 122-25).

It is clear that identification is an essential, if not the most critical, element of the postulated harms at the individual level, where much of the conventional privacy discussions focus. But even absent identification at the individual level, it may contribute to postulated harms at the collective or group level.

Group or collective harms may be illustrated thus. It is widely believed that there is greater consumption of adult or pornographic entertainment when conventions attended by large numbers of Christian Evangelicals are held at US hotels.[20] Whether true or false, this perception harms the collective image of Christian Evangelicals in the United States by showing them up as hypocrites.

To substantiate the above claim, it would not be necessary for hotels to release the video viewing records of individuals, an act that would violate the provisions of the US Video Privacy Protection Act of 1988. Instead, the hotels could simply provide the aggregate use records by title or category of videos together with the numbers of guests attending Evangelical and other conventions. With this information, it would be possible to observe the peaks and valleys of consumption of adult entertainment in hotels and their correlations with Evangelical and other conventions.

This is an example of a breach of collective or group privacy, as commonly understood. The simple aggregation of individual video rental records does not constitute the breach; it is the combination of that data with data identifying the group. The harm is connected to identification of the group.

It is critically important, however, to recognize the dangers associated with safeguarding "collective privacy" or "group privacy" of the type discussed above.

Rights are usually understood to belong to individuals, not to groups. The only group or collective right recognized in international law is that of peoples having the right of self-determination.[21] Even with this right, the value and operationalization of group rights are highly contested in the literature.[22]

---

[20] http://gospeldrivenchurch.blogspot.com/2011/03/what-you-do-in-your-hotel-room-gives.html. This site is sympathetic to Christians and hostile to adult entertainment.

[21] The United Nations, *International Covenant on Civil and Political Rights*, Article 1.

[22] Group rights, *Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/rights-group/

Furthermore, a prejudice against actions based on group attributes would pretty much put an end to efforts to improve the functioning of society in systematic, evidence-based ways. For example, it is routine to associate various characteristics or behaviors with persons living in geographical areas (e.g., in poverty mapping), by age group and gender and so on. It is considered desirable to "target" various policy measures to specific groups and indeed to improve the targeting by various means. Without group identification it will be impossible for modern societies to function. This is possibly the reason why safeguards against group identification do not currently exist and are not likely to exist in the future.

However, it has been reported that a book addressing questions the limitation of group rights to self-determining collectivities is about to be published by Luciano Floridi, Bart van der Sloot and Linnet Taylor of the University of Amsterdam.[23]

[23] The central argument is in Taylor (2015) and a response by Floridi (https://www.academia.edu/14389367/Open_Data_Data_Protection_and_Group_Privacy).

State of the art of technical solutions to masking identity

A common technique used by researchers to modify personal data such that the individual cannot be re-identified is called 'Anonymization.' The meaning of the term is broad but is more contentiously varied in its use. In its strongest sense anonymization attempts to prevent not just re-identification but to also prevent any information about specific subjects being learned. Such a strong anonymization technique reduces the utility of the data. Even simply just preventing re-identification by removal of personal identifiers can reduce the value of the analyses that can be performed. For example when one dataset includes multiple rows of data about a specific individual, removing the personal identifier makes each row akin to the record of a new individual.

More commonly in the case of using large datasets for developmental purposes, what is actually done is pseudonymization, whereby the personal identifiers in a dataset are replaced by unique identifiers disconnected to the real individual. Now the data points associated with the individual can be analyzed without connecting to the individual in realspace. The actual algorithms used to achieve pseudonymization may include, but are "not limited to preimage resistant hashes (e.g., one-way hashes) and encryption techniques in which the decryption key has been discarded" (United States Department of Justice, 2006, p. 12).

When the pseudonymized dataset is considered in isolation, the techniques may be made sufficiently robust to prevent re-identification using brute force techniques. But as the number and heterogeneity of datasets increases, re-identification becomes a possibility. This is because there may be 'pseudo-identifiers' such as aggregate demographic and/or geographic data even amongst pseudonymized data sets. These pseudo-identifiers can be potentially correlated with publicly available data (e.g., voter registration information) to identify the person. Hence current techniques for data anonymization (i.e., methods designed to strip data of PII) employed by computational social scientists, have been called into question.[24]

In one of the most publicized instances of anonymization reversal, Netflix released the viewing histories of 500,000 of its users without any PII in 2006, in an effort to improve its movie recommendation system. Narayanan & Shmatikov (2008) showed how these anonymized viewers could be re-identified if any of them had rated even a few movies on the International Movie Database (IMDB) website, where data is available publicly. Using this insight they were able to go further, even identifying identified user's religious and political leanings. Pseudonymized CDRs that are being used to produce insights of relevance for developmental policy have also been shown to be vulnerable to re-identification. Using CDRs for 1.5 million pseudonymized mobile subscribers covering a 15-month period, de Montjoye, Hidalgo, Verleysen and Blondel (2013) showed that up to 90 per cent of the subscribers could be 'identified' with just four data points, and 50 percent with just two data points. Although the actual identities of the users were unknown, the authors pointed out that the subscribers could in fact be completely re-identified by cross-referencing their results with other easily available data sources such as voter registration records.

[24] For further information on the range of anonymization techniques often utilized see El Emam (2013)

In light of these limitations, other techniques are being developed by computer scientists and by statistical scientists. These new techniques may be characterized under the broad heading of differential privacy, first introduced by Dwork (2006) and Dwork, et al. (2006), and built as extensions of other techniques such as K-anonymity and L-diversity that have been used widely when sharing health related data. K-anonymity and L-diversity attempt to reduce the chances for re-identification of individuals through the use of pseudo-identifiers. K-anonymity does that by trying to ensure that for each set of possible pseudo-identifiers, any re-identification attempt would return no less than K records, where K is usually greater than 1. L-diversity, works by trying to ensure that for every set of pseudo-identifiers that could be used in an re-identification attempt, there is more there is atleast L "well represented" values for certain attributes in the results of the original query that may be deemed to contain potentially confidential or sensitive information (e.g. income, age, etc.).[25] Differential privacy seeks to ensure that the results that are derived from a dataset are virtually the same whether a particular individual was in it or not. This is accomplished in principle by adding noise to the dataset in such a manner that it does not affect the overall statistical robustness of the results within a certain level of sensitivity.[26] By its very nature, this limits the scope of queries that can be conducted, which have to be aggregate queries. The distorting effects produced by differential privacy limits its use since too much noise may need to be added for most practical situations (Fienberg, Rinaldo, & Yang, 2010; Charest, 2012) especially when aggregate information is sought for groups of small sizes. Differential privacy is particular useful for large datasets with mainly categorical variables and where the dataset is sufficiently representative of the group for whom results are being sought. Census data are particularly suited for the use of differential privacy (Soria-Cormas & Drechsler, 2013). For now, differential privacy's mathematical constraints make it difficult to implement more broadly, especially for large semi-structured to unstructured datasets. Recent efforts in the application of differential privacy techniques to the analyses of MNBD (specifically CDRs) show promise. Mir et al. (2013) conducted experiments where the accuracy of results with the use of differential privacy techniques were high, suggesting that the mobility patterns of real metropolitan populations could be studied from CDR data whilst also preserving privacy.

The concerns regarding re-identification are valid, but they are also somewhat premature for developing countries given that overall levels of 'datafication' in developing economies are still quite low. The large majority of mobile phone connections in the developing world are prepaid, with minimal (if any) reliable associated registration information. Prompted by security concerns, governments are increasingly mandating the collection of registration information even for prepaid customers (GSMA, 2013). Even if these were mandatory often the registered prepaid user and the actual user may not be identical. SIM resellers may pre-register the SIMs they sell under their own name, or SIMs registered under the name of one family member may in fact be used by other members of the family as well. Sri Lankan operators have evidenced this mismatch. The same is also the case in many other developing countries.[27]

One can hope that new privacy-preserving techniques will be sufficiently advanced by the time developing economies become more 'datafied.' In the meantime, these new data sets can be

---

[25] For a more thorough treatment of these two techniques refer to Machanavajjhala & Reiter (2012).

[26] For a survey of differential privacy techniques see Ji, Lipton, & Elkan (2014)

[27] Based on interviews with operators in South Asia.

leveraged for public purposes under controlled situations backed by legal agreements and approval processes.

But it has to be acknowledged that even as the state of the art in privacy preserving techniques advances, privacy will need to be appreciated on a spectrum, with high privacy and low utility on one end, and low privacy and high utility on the other end. Where a particular analysis sits on this spectrum will at times necessitate non-technical solutions. For example legal agreements, approval processes and/or limited access-controls may all be needed for certain uses of the data. Mainstreaming the use of such data may then necessitate the need for some form of a priori privacy review that can look at privacy implications on a case-by-case basis. This could work in a fashion similar to extant ethics review boards that social scientists and academics frequency utilize.

Conclusions

Much of the discussion of the socio-economic implications of behavioral data has focused on the inclusion of more citizens and more aspects of their lives within the sphere of control enabled by pervasive data collection. This report examines marginalization or exclusion from the scope of data collection. It also examines frontier issues associated with behavioral big data, namely poverty/wealth mapping, including redlining, and the identification of regular and ad hoc congregations. In addition, it presents the state of the art on technical means used to mask PII in big data sets.

Marginalization has to be addressed as a special case of the problem of representivity. In the early days of big data, representivity was neglected to some extent. The problem is being paid increasing attention now by researchers. However, in the policy arena, there may be a tendency to act on insights that have been produced on the basis of available data. The solution is no different from that which is recommended to address the problem of representivity in general; explicitly address the absences. Especially in developing countries where datafication is rudimentary, the issue of representivity must be addressed in relation to the research questions being asked.

There is much interest in correlating socio-economic data with geographic locations in the form of poverty mapping. This necessarily involves wealth mapping too. Poverty mapping can help targeted delivery of services by government and relevant agencies. But a corollary is that knowledge of where people of wealth are concentrated may result in those areas being prioritized for delivery of certain forms of services as well and possibly the areas with concentrations of poverty. This would, in most countries, be unlawful or politically damaging if done by governments or monopolistic suppliers acting under authority of government.

In the case of competitive supply, some firms may choose to supply the wealthy areas while others may choose to concentrate on the poor areas. In general, this would not be unlawful and is natural under conditions of competitive supply.

Redlining, or the refusal to serve persons from specific geographical areas, is a phenomenon that has drawn the attention of policy makers in the United States. Historically, this has been done on the basis of crude correlations between location and ability to pay. In many cases this has also been correlated with ethnic identity.

Big data may enable forms of discrimination more precise than those associated with redlining. Algorithms may be used to mask unlawful forms of discrimination. They may also lead to more accurate identification of consumers with desirable or undesirable characteristics or propensities and end the crude and error-ridden forms of discrimination known as redlining. Indeed, data analytics may enable first-degree price discrimination, displacing traditional ways of pricing products.

MNDB and other forms of big data that yield insights on movement of people through time and space can allow the identification of regular and ad hoc congregations in specific locations.

Insights from pseudonymized historical data can be useful for deciding on locations of government and retail outlets and also for the pricing of outdoor advertising. Location-based advertising is of course a prime application. This can range from cell broadcasting which does not require piercing the collective shell and being able to differentiate between individual members of the congregation.

When it comes to analysis of real-time and non-anonymized data problems emerge. Participants in political protests may be identified and acted against, serious problem.

Issues of collective privacy apply to both aspects discussed above. While some degree of harm may occur, it is concluded that it is not advisable to extend privacy which is a valid concept at the individual level to the collective. This would negate most efforts to make efficient the delivery of public services.

In the case of technical methods of masking PII from individual data within big data sets, there is no easy solution though considerable advances have been made. In the case of developing countries, the current low levels of datafication offers safeguards. Until more sophisticated technical solutions are found the data sets should be used with non-technical safeguards such as legal agreements.

References

Arsu S. & Bilefsky D. (2014, March 21). In Turkey, Twitter Roars After Effort to Block It. *New York Times*. Retrieved from http://www.nytimes.com/2014/03/22/world/europe/turks-seek-to-challenge-twitter-ban.html

Beniger, James R. (1986). *The Control Revolution: Technological and Economic Origins of the Information Society*. Cambridge MA: Harvard U Press.

Borgesius, F; van Eechoud, M.; Gray, J. (2015). Open Data, Privacy, and Fair Information Principle, presentation to Berkeley, BCLT/BTLJ Symposium. https://www.law.berkeley.edu/files/Borgesius_van_Eechoud_Gray_IViR-Berkeley-slides_2015-04-08.pdf

Botta, F., Moat, H. S., Preis, T. (2015) Quantifying crowd size with mobile phone and Twitter data. R. Soc. open sci.2: 150162. http://dx.doi.org/10.1098/rsos.150162. Retrieved from http://rsos.royalsocietypublishing.org/content/royopensci/2/5/150162.full.pdf

Burns, John F. (2011 August 11). British Prime Minister Faces Questioning in House of Commons Over Rioting, *New York Times*, http://www.nytimes.com/2011/08/12/world/europe/12cameron.html?src=rec&recp=20#h[IslBti,1

Cate, F.; Cullen, P.; Mayer-Schonberger, V. (2013, December). *Data protection principles for the 21st century: Revising the OECD guidelines*. http://nova.ilsole24ore.com/wordpress/wp-content/uploads/2014/01/Data_Protection_Principles_for_the_21st_Century.pdf

Charest, A.-S. (2012). Empirical evaluation of statistical inference from differnetially-private contingency tables. In J. Domingo-Ferrer and I. Tinnirello, eds., *Privacy in Statistical Databases* (PSD2012), 257–272

David, T. (2013). Big Data from Cheap Phones. *Technology Review*, *116*(3), 50–54

De Montjoye, Y.-A.; Hidalgo, C.; Verleysen, M.; & Blondel, V. (2013), Unique in the Crowd: The privacy bounds of human mobility. *Scientific Reports*, 3, 1376. doi:10.1038/srep01376.

Dwork, C. (2006). Differential privacy. In *Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming* (ICALP), 1–12.

Dwork, C.; Mcsherry, F.; Nissim, K.; & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Theory of Cryptography Conference*, 265–284.

El Emam, K. (2013). *Guide to the de-identification of personal health information*. CRC Press, 2013.

Ewing, Jack (2015 October 4). Volkswagen Engine-Rigging Scheme Said to Have Begun in 2008, *New York Times*. http://www.nytimes.com/2015/10/05/business/engine-shortfall-pushed-volkswagen-to-evade-emissions-testing.html

Executive Office of the President of the United States. (2015, February). *Big Data and Differential Pricing*.

Fienberg, S.; Rinaldo, A.; & Yang, X. (2010). Differential privacy and the risk utility tradeoff for multi-dimensional contingency tables. In J. Domingo-Ferrer and E. Magkos, eds., *Privacy in Statistical Databases* (PSD2010), 187–199

Frias-Martinez, V.; Virseda-Jerez, J.; & Frias-Martinez, E. (2012). On the relation between socio-economic status and physical mobility. *Information Technology for Development*, *18*(2), 91–106. doi:10.1080/02681102.2011.630312

Goodman, L., (2015, March 7). The EMBERS Project Can Predict the Future With Twitter. *Newsweek*. Retrieved from http://www.newsweek.com/2015/03/20/embers-project-can-predict-future-twitter-312063.html

Gould-Wartofsky, M. (2015, May 5). From Ferguson to Baltimore, a 5-Step Guide to the Police Repression of Protest. *In These Times*. Retrieved from http://inthesetimes.com/article/17909/ferguson_baltimore_police_repression

GSMA (2013). The Mandatory Registration of Prepaid SIM Card Users. Retrieved from: http://www.gsma.com/publicpolicy/wp-content/uploads/2013/11/GSMA_White-Paper_Mandatory-Registration-of-Prepaid-SIM-Users_32pgWEBv3.pdf

Gutierrez, T.; Krings, G.; & Blondel, V. (2013). Evaluating socio-economic state of a country analyzing airtime credit and mobile phone datasets, 1–6. Retrieved from http://arxiv.org/abs/1309.4496

Hannak, A.: Soeller, G.; Lazer, D.; Mislove, A.; Wilson, C. (2014). Measuring Price Discrimination and Steering on E-commerce Web Sites. http://www.ccs.neu.edu/home/cbw/pdf/imc151-hannak.pdf

Ji, Z., Lipton, Z. C., & Elkan, C. (2014). Differential Privacy and Machine Learning: a Survey and Review, 1–30. Learning; Cryptography and Security; Databases. Retrieved from http://arxiv.org/abs/1412.7584

Kahneman, D; Knetsch, J. and Thaler, R (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review*, 76(4).

Kallus, N. (2014). Predicting Crowd behavior with big public data. WWW'14 Companion, April 7–11, 2014, Seoul, Korea ACM 978-1-4503-2745-9/14/04. Retrieved from http://arxiv.org/pdf/1402.2308v1.pdf

Lerman, J. (2013). Big data and its exclusions. *Stanford Law Review Online*, 66. Retrieved from http://www.stanfordlawreview.org/online/privacy-and-big-data/big-data-and-its-exclusions

Lopez, T. (2014, January 24) How Did Ukraine's Government Text Threats to Kiev's EuroMaidan Protesters?. *Slate.* Retrieved from http://www.slate.com/blogs/future_tense/2014/01/24/ukraine_texting_euromaidan_protesters_kiev_demonstrators_receive_threats.html

Machanavajjhala, A., & Reiter, J. P. (2012). Big privacy. *XRDS: Crossroads, The ACM Magazine for Students*. doi:10.1145/2331042.2331051.

Mayer-Schonberger, V.; Cukier, K. (2013). *Big data*. London: John Murray.

Miller M.; Ginnis S.; Stobart, R.; Krasodomski-Jones, A.; Clemence, M. (2015). *The road to representivity; a Demos and Ipsos MORI report on sociological research using Twitte*r. Retrieved from https://www.ipsos-mori.com/Assets/Docs/Publications/ipsos-mori-demos-road-to-representivity.pdf

Mir, D. J., Isaacman, S., Caceres, R., Martonosi, M., & Wright, R. N. (2013, October). Dp-where: Differentially private modeling of human mobility. In *Big Data, 2013 IEEE International Conference on* (pp. 580-588). IEEE.

Mundie, C (2014). Privacy pragmatism. *Foreign Affairs*. March/April. http://www.foreignaffairs.com/articles/140741/craig-mundie/privacy-pragmatism

Narayanan, A. and Shmatikov, V. (2008), Robust de-anonymization of large sparse datasets. *2008 IEEE Symposium on Security and Privacy* (sp 2008) (pp. 111–125). IEEE. doi:10.1109/SP.2008.33

Parkinson, J. (2013, June 3). Amid Turkey Unrest, Social Media Becomes a Battleground. *Wall Street Journal*. Retrieved from http://blogs.wsj.com/middleeast/2013/06/03/amid-turkey-unrest-social-media-becomes-a-battleground/

Pasquale, Frank (2015*). The black box society: The secret algorithms that control money and information*. Cambridge MA: Harvard University Press.

Podesta J.; Pritzker, P.; Moniz, E.; Holdren, J.; Zients J. (2014, May) *Big Data: Seizing Opportunities, Preserving Values*. Executive Office of the President. Retrieved from https://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf

Samarajiva, R. (1994). Privacy in electronic public space, *Canadian Journal of Communication*, 19(1): 87-99.

Samarajiva, R. (2014, February 25). Big data in developing v. developed countries. http://lirneasia.net/2014/02/big-data-in-developing-v-developed-countries/

Samarajiva, R.; Lokanathan, S.; Madhawa, K.; Kriendler, G., & Maldeniya, D. (2015). "Big data to improve urban planning," *Economic and Political Weekly*, Vol L. No. 22, May 30: 42-48) http://www.epw.in/review-urban-affairs/big-data-improve-urban-planning.html

Savage, C. (2013 October 18). NSA plan to log calls is renewed by court. *New York Times*, http://www.nytimes.com/2013/10/19/us/nsa-plan-to-log-calls-is-renewed-by-court.html?module=Search&mabReward=relbias%3Aw

Shiller, B. R. (2014, January 30). First Degree price discrimination using Big Data. *Brandeis University*. Retrieved from http://benjaminshiller.com/images/First_Degree_PD_Using_Big_Data_Jan_27,_2014.pdf

Solove, D.J. (2008). *Understanding privacy*. Cambridge MA: Harvard University Press.

Soria-Cormas, J. & Drechsler, J. (2013). Evaluating the potential of differential privacy mechanisms for census data. In *UNECE Conference of European Statisticians*.

Taylor, L. (2015). No place to hide? The ethics and analytics of tracking mobility using mobile phone data, *Environment and Planning D: Society and Space*.

Tufekci, Zeynep (2014, July). Engineering the public: Big data, surveillance and computational politics. *First Monday*, 19(7)

United States Department of Justice. (2006). *Privacy Technology Focus Group Report*

Valentino-Devries, J.; Singer-Vine, J. (2012, December 24). Websites Vary Prices, Deals Based on Users' Information. *Wall Street Journal*. Retrieved from http://www.wsj.com/articles/SB10001424127887323777204578189391813881534

Walker, S.; Grytsenko, O. (2014, Jan 21). Text messages warn Ukraine protesters they are 'participants in mass riot. *The Guardian.* Retrieved from http://www.theguardian.com/world/2014/jan/21/ukraine-unrest-text-messages-protesters-mass-riot