

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/317139233>

The Tyranny of Data? The Bright and Dark Sides of Data-Driven Decision-Making for Social Good

Chapter · May 2017

DOI: 10.1007/978-3-319-54024-5_1

CITATIONS

7

READS

172

5 authors, including:



Bruno Lepri

Fondazione Bruno Kessler

113 PUBLICATIONS 1,761 CITATIONS

SEE PROFILE



Jacopo Staiano

Università degli Studi di Trento

33 PUBLICATIONS 732 CITATIONS

SEE PROFILE



Emmanuel Francis Letouzé

Massachusetts Institute of Technology

9 PUBLICATIONS 57 CITATIONS

SEE PROFILE



Nuria Oliver

Vodafone

183 PUBLICATIONS 8,427 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Modeling Dominance Effects on Nonverbal Behaviour with Granger Causality [View project](#)



Psychographics [View project](#)

The Tyranny of Data? The Bright and Dark Sides of Data-Driven Decision-Making for Social Good

Bruno Lepri, Jacopo Staiano, David Sangokoya, Emmanuel Letouzé and
Nuria Oliver

Abstract The unprecedented availability of large-scale human behavioral data is profoundly changing the world we live in. Researchers, companies, governments, financial institutions, non-governmental organizations and also citizen groups are actively experimenting, innovating and adapting algorithmic decision-making tools to understand global patterns of human behavior and provide decision support to tackle problems of societal importance. In this chapter, we focus our attention on *social good decision-making algorithms*, that is algorithms strongly influencing decision-making and resource optimization of public goods, such as public health, safety, access to finance and fair employment. Through an analysis of specific use cases and approaches, we highlight both the positive opportunities that are created through data-driven algorithmic decision-making, and the potential negative consequences that practitioners should be aware of and address in order to truly realize the potential of this emergent field. We elaborate on the need for these algorithms to provide transparency and accountability, preserve privacy and be tested and evaluated *in context*, by means of living lab approaches involving citizens. Finally, we turn to the requirements which would make it possible to leverage the predictive power of data-driven human behavior analysis while ensuring transparency, accountability, and civic participation.

Bruno Lepri
Fondazione Bruno Kessler e-mail: lepri@fbk.eu

Jacopo Staiano
Fortia Financial Solutions e-mail: jacopo.staiano@fortia.fr

David Sangokoya
Data-Pop Alliance e-mail: dsangokoya@datapopalliance.org

Emmanuel Letouzé
Data-Pop Alliance and MIT Media Lab e-mail: eletouze@mit.edu

Nuria Oliver
Data-Pop Alliance e-mail: nuria@alum.mit.edu

1 Introduction

The world is experiencing an unprecedented transition where human behavioral data has evolved from being a scarce resource to being a massive and real-time stream. This availability of large-scale data is profoundly changing the world we live in and has led to the emergence of a new discipline called *computational social science* [45]; finance, economics, marketing, public health, medicine, biology, politics, urban science and journalism, to name a few, have all been disrupted to some degree by this trend [41].

Moreover, the automated analysis of anonymized and aggregated large-scale human behavioral data offers new possibilities to understand global patterns of human behavior and to help decision makers tackle problems of societal importance [45], such as monitoring socio-economic deprivation [8, 75, 76, 88] and crime [11, 10, 84, 85, 90], mapping the propagation of diseases [37, 94], or understanding the impact of natural disasters [55, 62, 97]. Thus, researchers, companies, governments, financial institutions, non-governmental organizations and also citizen groups are actively experimenting, innovating and adapting algorithmic decision-making tools, often relying on the analysis of personal information.

However, researchers from different disciplinary backgrounds have identified a range of social, ethical and legal issues surrounding data-driven decision-making, including privacy and security [19, 22, 23, 56], transparency and accountability [18, 61, 99, 100], and bias and discrimination [3, 79]. For example, Barocas and Selbst [3] point out that the use of data-driven decision making processes can result in disproportionate adverse outcomes for disadvantaged groups, in ways that look like discrimination. Algorithmic decisions can reproduce patterns of discrimination, due to decision makers' prejudices [60], or reflect the biases present in the society [60]. In 2014, the White House released a report, titled "Big Data: Seizing opportunities, preserving values" [65] that highlights the discriminatory potential of big data, including how it could undermine longstanding civil rights protections governing the use of personal information for credit, health, safety, employment, etc. For example, data-driven decisions about applicants for jobs, schools or credit may be affected by hidden biases that tend to flag individuals from particular demographic groups as unfavorable for such opportunities. Such outcomes can be self-reinforcing, since systematically reducing individuals' access to credit, employment and educational opportunities may worsen their situation, which can play against them in future applications.

In this chapter, we focus our attention on *social good algorithms*, that is algorithms strongly influencing decision-making and resource optimization of public goods, such as public health, safety, access to finance and fair employment. These algorithms are of particular interest given the magnitude of their impact on quality of life and the risks associated with the information asymmetry surrounding their governance.

In a recent book, William Easterly evaluates how global economic development and poverty alleviation projects have been governed by a “tyranny of experts” – in this case, aid agencies, economists, think tanks and other analysts – who consistently favor top-down, technocratic governance approaches at the expense of the individual rights of citizens [28]. Easterly details how these experts reduce multidimensional social phenomena such as poverty or justice into a set of technical solutions that do not take into account either the political systems in which they operate or the rights of intended beneficiaries. Take for example the displacement of farmers in the Mubende district of Uganda: as a direct result of a World Bank project intended to raise the region’s income by converting land to higher value uses, farmers in this district were forcibly removed from their homes by government soldiers in order to prepare for a British company to plant trees in the area [28]. Easterly underlines the cyclic nature of this tyranny: technocratic justifications for specific interventions are considered objective; intended beneficiaries are unaware of the opaque, black box decision-making involved in these resource optimization interventions; and experts (and the coercive powers which employ them) act with impunity and without redress.

If we turn to the use, governance and deployment of big data approaches in the public sector, we can draw several parallels towards what we refer to as the “tyranny of data”, that is the adoption of data-driven decision-making under the technocratic and top-down approaches highlighted by Easterly [28]. We elaborate on the need for *social good decision-making algorithms* to provide transparency and accountability, to only use personal information – owned and controlled by individuals – with explicit consent, to ensure that privacy is preserved when data is analyzed in aggregated and anonymized form, and to be tested and evaluated *in context*, that is by means of living lab approaches involving citizens. In our view, these characteristics are crucial for fair data-driven decision-making as well as for citizen engagement and participation.

In the rest of this chapter, we provide the readers with a compendium of the issues arising from current big data approaches, with a particular focus on specific use cases that have been carried out to date, including urban crime prediction [10], inferring socioeconomic status of countries and individuals [8, 49, 76], mapping the propagation of diseases [37, 94] and modeling individuals’ mental health [9, 20, 47]. Furthermore, we highlight factors of risk (*e.g.* privacy violations, lack of transparency and discrimination) that might arise when decisions potentially impacting the daily lives of people are heavily rooted in the outcomes of black-box data-driven predictive models. Finally, we turn to the requirements which would make it possible to leverage the predictive power of data-driven human behavior analysis while ensuring transparency, accountability, and civic participation.

2 The rise of data-driven decision-making for social good

The unprecedented stream of large-scale, human behavioral data has been described as a “tidal wave” of opportunities to both predict and act upon the analysis of the petabytes of digital signals and traces of human actions and interactions. With such massive streams of relevant data to mine and train algorithms with, as well as increased analytical and technical capacities, it is of no surprise that companies and public sector actors are turning to machine learning-based algorithms to tackle complex problems at the limits of human decision-making [36, 96]. The history of human decision-making – particularly when it comes to questions of power in resource allocation, fairness, justice, and other public goods – is wrought with innumerable examples of extreme bias, leading towards corrupt, inefficient or unjust processes and outcomes [2, 34, 70, 87]. In short, human decision-making has shown significant limitations and the turn towards data-driven algorithms reflects a search for objectivity, evidence-based decision-making, and a better understanding of our resources and behaviors.

Diakopoulos [27] characterizes the function and power of algorithms in four broad categories: 1) *classification*, the categorization of information into separate “classes”, based on its features; 2) *prioritization*, the denotation of emphasis and rank on particular information or results at the expense of others based on a pre-defined set of criteria; 3) *association*, the determination of correlated relationships between entities; and 4) *filtering*, the inclusion or exclusion of information based on pre-determined criteria.

Table 1 provides examples of types of algorithms across these categories.

Table 1 Algorithmic function and examples, adapted from Diakopoulos [27] and Latzer *et al.* [44]

Function	Type	Examples
<i>Prioritization</i>	General and search engines, meta search engines, semantic search engines, questions & answers services	Google, Bing, Baidu; image search; social media; Quora; Ask.com
<i>Classification</i>	Reputation systems, news scoring, credit scoring, social scoring	Ebay, Uber, Airbnb; Reddit, Digg; CreditKarma; Klout
<i>Association</i>	Predicting developments and trends	ScoreAhit, Music Xray, Google Flu Trends
<i>Filtering</i>	Spam filters, child protection filters, recommender systems, news aggregators	Norton; Net Nanny; Spotify, Netflix; Facebook Newsfeed

This chapter places emphasis on what we call social good algorithms – algorithms strongly influencing decision-making and resource optimization for

public goods. These algorithms are designed to analyze massive amounts of human behavioral data from various sources and then, based on pre-determined criteria, select the information most relevant to their intended purpose. While resource allocation and decision optimization over limited resources remain common features of the public sector, the use of social good algorithms brings to a new level the amount of human behavioral data that public sector actors can access, the capacities with which they can analyze this information and deliver results, and the communities of experts and common people who hold these results to be objective. The ability of these algorithms to identify, select and determine information of relevance beyond the scope of human decision-making creates a new kind of decision optimization facilitated by both the design of the algorithms and the data from which they are based. However, as discussed later in the chapter, this new process is often opaque and assumes a level of impartiality that is not always accurate. It also creates information asymmetry and lack of transparency between actors using these algorithms and the intended beneficiaries whose data is being used.

In the following sub-sections, we assess the nature, function and impact of the use of social good algorithms in three key areas: criminal behavior dynamics and predictive policing; socio-economic deprivation and financial inclusion; and public health.

2.1 Criminal behavior dynamics and predictive policing

Researchers have turned their attention to the automatic analysis of criminal behavior dynamics both from a *people*- and a *place*-centric perspectives. The people-centric perspective has mostly been used for individual or collective criminal profiling [67, 72, 91]. For example, Wang *et al.* [91] proposed a machine learning approach, called Series Finder, to the problem of detecting specific patterns in crimes that are committed by the same offender or group of offenders.

In 2008, the criminologist David Weisburd proposed a shift from a people-centric paradigm of police practices to a place-centric one [93], thus focusing on geographical topology and micro-structures rather than on criminal profiling. An example of a place-centric perspective is the detection, analysis, and interpretation of crime hotspots [16, 29, 53]. Along these lines, a novel application of quantitative tools from mathematics, physics and signal processing has been proposed by Toole *et al.* [84] to analyse spatial and temporal patterns in criminal offense records. Their analyses of crime data from 1991 to 1999 for the American city of Philadelphia indicated the existence of multi-scale complex relationships in space and time. Further, over the last few years, aggregated and anonymized mobile phone data has opened new possibilities to study city dynamics with unprecedented temporal and spatial granular-

ities [7]. Recent work has used this type of data to predict crime hotspots through machine-learning algorithms [10, 11, 85].

More recently, these *predictive policing* approaches [64] are moving from the academic realm (universities and research centers) to police departments. In Chicago, police officers are paying particular attention to those individuals flagged, through risk analysis techniques, as most likely to be involved in future violence. In Santa Cruz, California, the police have reported a dramatic reduction in burglaries after adopting algorithms that predict where new burglaries are likely to occur. In Charlotte, North Carolina, the police department has generated a map of high-risk areas that are likely to be hit by crime. The Police Departments of Los Angeles, Atlanta and more than 50 other cities in the US are using PredPol, an algorithm that generates 500 by 500 square foot predictive boxes on maps, indicating areas where crime is most likely to occur. Similar approaches have also been implemented in Brasil, the UK and the Netherlands. Overall, four main predictive policing approaches are currently being used: (i) methods to forecast places and times with an increased risk of crime [32], (ii) methods to detect offenders and flag individuals at risk of offending in the future [64], (iii) methods to identify perpetrators [64], and (iv) methods to identify groups or, in some cases, individuals who are likely to become the victims of crime [64].

2.2 *Socio-economic deprivation and financial inclusion*

Being able to accurately measure and monitor key sociodemographic and economic indicators is critical to design and implement public policies [68]. For example, the geographic distribution of poverty and wealth is used by governments to make decisions about how to allocate scarce resources and provides a foundation for the study of the determinants of economic growth [33, 43]. The quantity and quality of economic data available have significantly improved in recent years. However, the scarcity of reliable key measures in developing countries represents a major challenge to researchers and policy-makers¹, thus hampering efforts to target interventions effectively to areas of greatest need (*e.g.* African countries) [26, 40]. Recently, several researchers have started to use mobile phone data [8, 49, 76], social media [88] and satellite imagery [39] to infer the poverty and wealth of individual subscribers, as well as to create high-resolution maps of the geographic distribution of wealth and deprivation.

The use of novel sources of behavioral data and algorithmic decision-making processes is also playing a growing role in the area of financial services, for example credit scoring. Credit scoring is a widely used tool in the financial sector to compute the risks of lending to potential credit customers. Providing

¹ <http://www.undatarevolution.org/report/>

information about the ability of customers to pay back their debts or conversely to default, credit scores have become a key variable to build financial models of customers. Thus, as lenders have moved from traditional interview-based decisions to data-driven models to assess credit risk, consumer lending and credit scoring have become increasingly sophisticated. Automated credit scoring has become a standard input into the pricing of mortgages, auto loans, and unsecured credit. However, this approach is mainly based on the past financial history of customers (people or businesses) [81], and thus not adequate to provide credit access to people or businesses when no financial history is available. Therefore, researchers and companies are investigating novel sources of data to replace or to improve traditional credit scores, potentially opening credit access to individuals or businesses that traditionally have had poor or no access to mainstream financial services –*e.g.* people who are unbanked or underbanked, new immigrants, graduating students, etc. Researchers have leveraged mobility patterns from credit card transactions [73] and mobility and communication patterns from mobile phones to automatically build user models of spending behavior [74] and propensity to credit defaults [71, 73]. The use of mobile phone, social media, and browsing data for financial risk assessment has also attracted the attention of several entrepreneurial efforts, such as Cignifi², Lenddo³, InVenture⁴, and ZestFinance⁵.

2.3 Public health

The characterization of individuals and entire populations' mobility is of paramount importance for public health [57]: for example, it is key to predict the spatial and temporal risk of diseases [35, 82, 94], to quantify exposure to air pollution [48], to understand human migrations after natural disasters or emergency situations [4, 50], etc. The traditional approach has been based on household surveys and information provided from census data. These methods suffer from recall bias and limitations in the size of the population sample, mainly due to excessive costs in the acquisition of the data. Moreover, survey or census data provide a snapshot of the population dynamics at a given moment in time. However, it is fundamental to monitor mobility patterns in a continuous manner, in particular during emergencies in order to support decision making or assess the impact of government measures.

Tizzoni *et al.* [82] and Wesolowski *et al.* [95] have compared traditional mobility surveys with the information provided by mobile phone data (Call

² <http://cignifi.com/>

³ <https://www.lenddo.com/>

⁴ <http://tala.co/>

⁵ <https://www.zestfinance.com/>

Detail Records or CDRs), specifically to model the spread of diseases. The findings of these works recommend the use of mobile phone data, by themselves or in combination with traditional sources, in particular in low-income economies where the availability of surveys is highly limited.

Another important area of opportunity within public health is mental health. Mental health problems are recognized to be a major public health issue⁶. However, the traditional model of episodic care is suboptimal to prevent mental health outcomes and improve chronic disease outcomes. In order to assess human behavior in the context of mental wellbeing, the standard clinical practice relies on periodic self-reports that suffer from subjectivity and memory biases, and are likely influenced by the current mood state. Moreover, individuals with mental conditions typically visit doctors when the crisis has already happened and thus report limited information about precursors useful to prevent the crisis onset. These novel sources of behavioral data yield the possibility of monitoring mental health-related behaviors and symptoms outside of clinical settings and without having to depend on self-reported information [52]. For example, several studies have shown that behavioral data collected through mobile phones and social media can be exploited to recognize bipolar disorders [20, 30, 59], mood [47], personality [25, 46] and stress [9].

Table 2 summarizes the main points emerging from the literature reviewed in this section.

Table 2 Summary table for the literature discussed in Section 2.

Key Area	Problems Tackled	References
<i>Predictive Policing</i>	Criminal behavior profiling	[67, 72, 91]
	Crime hotspot prediction	[10, 11, 32, 85]
	Perpetrator(s)/victim(s) identification	[64]
<i>Finance & Economy</i>	Wealth & deprivation mapping	[8, 49, 39, 76, 88]
	Spending behavior profiling	[74]
	Credit scoring	[71, 73]
<i>Public Health</i>	Epidemiologic studies	[35, 82, 94]
	Environment and emergency mapping	[4, 48, 50]
	Mental Health	[9, 20, 25, 30, 46, 47, 52, 59]

3 The dark side of data-driven decision-making for social good

The potential positive impact of big data and machine learning-based approaches to decision-making is huge. However, several researchers and ex-

⁶ http://www.who.int/topics/mental_health/en/

perts [3, 19, 61, 79, 86] have underlined what we refer to as *the dark side* of data-driven decision-making, including violations of privacy, information asymmetry, lack of transparency, discrimination and social exclusion. In this section we turn our attention to these elements before outlining three key requirements that would be necessary in order to realize the positive impact, while minimizing the potential negative consequences of data-driven decision-making in the context of social good.

3.1 Computational violations of privacy

Reports and studies [66] have focused on the misuse of personal data disclosed by users and on the aggregation of data from different sources by entities playing as data brokers with direct implications in privacy. An often overlooked element is that the computational developments coupled with the availability of novel sources of behavioral data (*e.g.* social media data, mobile phone data, etc.) now allow inferences about private information that may never have been disclosed. This element is essential to understand the issues raised by these algorithmic approaches.

A recent study by Kosinski *et al.* [42] combined data on Facebook “Likes” and limited survey information to accurately predict a male user’s sexual orientation, ethnic origin, religious and political preferences, as well as alcohol, drugs, and cigarettes use. Moreover, Twitter data has recently been used to identify people with a high likelihood of falling into depression before the onset of the clinical symptoms [20].

It has also been shown that, despite the algorithmic advancements in anonymizing data, it is feasible to infer identities from anonymized human behavioral data, particularly when combined with information derived from additional sources. For example, Zang *et al.* [98] have reported that if home and work addresses were available for some users, up to 35% of users of the mobile network could be de-identified just using the two most visited towers, likely to be related to their home and work location. More recently, de Montjoye *et al.* [22, 23] have demonstrated how unique mobility and shopping behaviors are for each individual. Specifically, they have shown that four spatio-temporal points are enough to uniquely identify 95% of people in a mobile phone database of 1.5M people and to identify 90% of people in a credit card database of 1M people.

3.2 Information asymmetry and lack of transparency

Both governments and companies use data-driven algorithms for decision making and optimization. Thus, accountability in government and corporate

use of such decision making tools is fundamental in both validating their utility toward the public interest as well as redressing harms generated by these algorithms.

However, the ability to accumulate and manipulate behavioral data about customers and citizens on an unprecedented scale may give big companies and intrusive/authoritarian governments powerful means to manipulate segments of the population through targeted marketing efforts and social control strategies. In particular, we might witness an *information asymmetry* situation where a powerful few have access and use knowledge that the majority do not have access to, thus leading to an –or exacerbating the existing– asymmetry of power between the state or the big companies on one side and the people on the other side [1]. In addition, the nature and the use of various data-driven algorithms for social good, as well as the lack of computational or data literacy among citizens, makes algorithmic transparency difficult to generalize and accountability difficult to assess [61].

Burrell [12] has provided a useful framework to characterize three different types of opacity in algorithmic decision-making: (1) *intentional opacity*, whose objective is the protection of the intellectual property of the inventors of the algorithms. This type of opacity could be mitigated with legislation that would force decision-makers towards the use of open source systems. The new General Data Protection Regulations (GDPR) in the EU with a “right to an explanation” starting in 2018 is an example of such legislation⁷. However, there are clear corporate and governmental interests in favor of intentional opacity which make it difficult to eliminate this type of opacity; (2) *illiterate opacity*, due to the fact that the vast majority of people lack the technical skills to understand the underpinnings of algorithms and machine learning models built from data. This kind of opacity might be attenuated with stronger education programs in computational thinking and by enabling that independent experts advice those affected by algorithm decision-making; and (3) *intrinsic opacity*, which arises by the nature of certain machine learning methods that are difficult to interpret (*e.g.* deep learning models). This opacity is well known in the machine learning community (usually referred to as the *interpretability problem*). The main approach to combat this type of opacity requires using alternative machine learning models that are easy to interpret by humans, despite the fact that they might yield lower accuracy than black-box non-interpretable models.

Fortunately, there is increasing awareness of the importance of reducing or eliminating the opacity of data-driven algorithmic decision-making systems. There are a number of research efforts and initiatives in this direction, including the Data Transparency Lab⁸ which is a “community of technolo-

⁷ Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) <http://eur-lex.europa.eu/eli/reg/2016/679/oj>

⁸ <http://www.datatransparencylab.org/>

gists, researchers, policymakers and industry representatives working to advance online personal data transparency through research and design”, and the DARPA Explainable Artificial Intelligence (XAI) project⁹. A tutorial on the subject has been held at the 2016 ACM Knowledge and Data Discovery conference [38]. Researchers from New York University’s Information Law Institute, such as Helen Nissenbaum and Solon Barocas, and Microsoft Research, such as Kate Crawford and Tarleton Gillespie, have held several workshops and conferences during the past few years on the ethical and legal challenges related to algorithmic governance and decision-making.¹⁰ A nominee for the National Book Award, Cathy O’Neil’s book, “Weapons of Math Destruction,” details several case studies on harms and risks to public accountability associated with big data-driven algorithmic decision-making, particularly in the areas of criminal justice and education [58]. Recently, in partnership with Microsoft Research and others, the White House Office of Science and Technology Policy has co-hosted several public symposiums on the impacts and challenges of algorithms and artificial intelligence, specifically in social inequality, labor, healthcare and ethics.¹¹

3.3 Social exclusion and discrimination

From a legal perspective, Tobler [83] argued that discrimination derives from “the application of different rules or practices to comparable situations, or of the same rule or practice to different situations”. In a recent paper, Barocas and Selbst [3] elaborate that discrimination may be an artifact of the data collection and analysis process itself; more specifically, even with the best intentions, data-driven algorithmic decision-making can lead to discriminatory practices and outcomes. Algorithmic decision procedures can reproduce existing patterns of discrimination, inherit the prejudice of prior decision makers, or simply reflect the widespread biases that persist in society [19]. It can even have the perverse result of exacerbating existing inequalities by suggesting that historically disadvantaged groups actually deserve less favorable treatment [58].

Discrimination from algorithms can occur for several reasons. First, input data into algorithmic decisions may be poorly weighted, leading to *disparate impact*; for example, as a form of *indirect discrimination*, overemphasis of zip code within predictive policing algorithms can lead to the association of low-income African-American neighborhoods with areas of crime and as a result, the application of specific targeting based on group membership [17]. Second, discrimination can occur from the decision to use an algorithm itself.

⁹ <http://www.darpa.mil/program/explainable-artificial-intelligence>

¹⁰ <http://www.law.nyu.edu/centers/ili/algorithmsconference>

¹¹ <https://www.whitehouse.gov/blog/2016/05/03/preparing-future-artificial-intelligence>

Categorization – through algorithmic classification, prioritization, association and filtering – can be considered as a form of *direct discrimination*, whereby algorithms are used for disparate treatment [27]. Third, algorithms can lead to discrimination as a result of the misuse of certain models in different contexts [14]. Fourth, in a form of feedback loop, biased training data can be used both as evidence for the use of algorithms and as proof of their effectiveness [14].

The use of algorithmic data-driven decision processes may also result in individuals mistakenly being denied opportunities based not on their own action but on the actions of others with whom they share some characteristics. For example, some credit card companies have lowered a customer’s credit limit, not based on the customer’s payment history, but rather based on analysis of other customers with a poor repayment history that had shopped at the same establishments where the customer had shopped [66].

Indeed, we find increasing evidence of detrimental impact already taking place in current non-algorithmic approaches to credit scoring and generally to backgrounds checks. The latter have been widely used in recent years in several contexts: it is common to agree to be subjected to background checks when applying for a job, to lease a new apartment, etc. In fact, hundreds of thousands of people have unknowingly seen themselves adversely affected on existential matters such as job opportunities and housing availability due to simple but common mistakes (for instance, misidentification) in the procedures used by external companies to perform background checks¹². It is worth noticing that the trivial procedural mistakes causing such adverse outcomes are bound to disappear once fully replaced with data-driven methodologies. Alas, this also means that should such methodologies not be transparent in their inner workings, the effects are likely to stay though with different roots. Further, the effort required to identify the causes of unfair and discriminative outcomes can be expected to be exponentially larger, as exponentially more complex will be the black-box models employed to assist in the decision-making process. This scenario highlights particularly well the need for machine learning models featuring transparency and accountability: adopting black-box approaches in scenarios where the lives of people would be seriously affected by a machine-driven decision could lead to forms of *algorithmic stigma*¹³, a particularly creepy scenario considering that those stigmatized might never become aware of being so, and the stigmatizer will be an unaccountable machine. Recent advances in neural network-based (deep learning) models are yielding unprecedented accuracies in a variety of fields. However, such models tend to be difficult – if not impossible – to interpret, as

¹² See, for instance, <http://www.chicagotribune.com/business/ct-background-check-penalties-1030-biz-20151029-story.html>

¹³ As a social phenomenon, the concept of stigma has received significant attention by sociologists, who under different frames highlighted and categorized the various factors leading individuals or groups to be discriminated by society, the countermoves often adopted by the stigmatized, and analyzed dynamics of reactions and evolution of stigma. We refer the interested reader to the review provided by Major and O’Brian [51].

previously explained. In this chapter, we highlight the need for data-driven machine learning models that are interpretable by humans when such models are going to be used to make decisions that affect individuals or groups of individuals.

4 Requirements for positive disruption of data-driven policies

As noted in the previous sections, both governments and companies are increasingly using data-driven algorithms for decision support and resource optimization. In the context of social good, accountability in the use of such powerful decision support tools is fundamental in both validating their utility toward the public interest as well as redressing corrupt or unjust harms generated by these algorithms. Several scholars have emphasized elements of what we refer to as the dark side of data-driven policies for social good, including violations of individual and group privacy, information asymmetry, lack of transparency, social exclusion and discrimination. Arguments against the use of social good algorithms typically call into question the use of machines in decision support and the need to protect the role of human decision-making.

However, therein lies a huge potential and imperative for leveraging large scale human behavioral data to design and implement policies that would help improve the lives of millions of people. Recent debates have focused on characterizing data-driven policies as either “good” or “bad” for society. We focus instead on the potential of data-driven policies to lead to positive disruption, such that they reinforce and enable the powerful functions of algorithms as tools generating value while minimizing their dark side.

In this section, we present key *human-centric* requirements for positive disruption, including a fundamental renegotiation of user-centric data ownership and management, the development of tools and participatory infrastructures towards increased algorithmic transparency and accountability, and the creation of living labs for experimenting and co-creating data-driven policies. We place humans at the center of our discussion as humans are ultimately both the actors and the subjects of the decisions made via algorithmic means. If we are able to ensure that these requirements are met, we should be able to realize the positive potential of data-driven algorithmic decision-making while minimizing the risks and possible negative unintended consequences.

4.1 *User-centric data ownership and management*

A big question on the table for policy-makers, researchers, and intellectuals is: *how do we unlock the value of human behavioral data while preserving*

the fundamental right to privacy? This question implicitly recognizes the risks, in terms not only of possible abuses but also of a “missed chance for innovation”, inherent to the current paradigm: the dominant *siloed* approach to data collection, management, and exploitation, precludes participation to a wide range of actors, most notably to the very producers of personal data (*i.e.* the users).

On this matter, new user-centric models for personal data management have been proposed, in order to empower individuals with more control of their own data’s life-cycle [63]. To this end, researchers and companies are developing repositories which implement medium-grained access control to different kinds of personally identifiable information (PII), such as passwords, social security numbers and health data [92], location [24] and personal data collected by means of smartphones or connected devices [24]. A pillar of these approaches is represented by a *Personal Data Eco-system*, composed by secure vaults of personal data whose owners are granted full control of.

Along this line, an interesting example is the Enigma platform [101] that leverages the recent technological trend of decentralization: advances in the fields of cryptography and decentralized computer networks have resulted in the emergence of a novel technology – known as the *blockchain* – which has the potential to reduce the role of one of the most important actors in our society: the middle man [5, 21]. By allowing people to transfer a unique piece of digital property or data to others, in a safe, secure, and immutable way, this technology can create digital currencies (*e.g.* bitcoin) that are not backed by any governmental body [54]; self-enforcing digital contracts, called smart contracts, whose execution does not require any human intervention (*e.g.* Ethereum) [80]; and decentralized marketplaces that aim to operate free from regulations [21]. Hence, Enigma tackles the challenge of providing a secure and trustworthy mechanism for the exchange of goods in a personal data market. To illustrate how the platform works, consider the following example: a group of data analysts of an insurance company wishes to test a model that leverages people’s mobile phone data. Instead of sharing their raw data with the data analysts in the insurance company, the users can securely store their data in Enigma, and only provide the data analysts with a permission to execute their study. The data analysts are thus able to execute their code and obtain the results, but nothing else. In the process, the users are compensated for having given access to their data and the computers in the network are paid for their computing resources [78].

4.2 Algorithmic transparency and accountability

The deployment of a machine learning model entails a degree of *trust* on how satisfactory its performance in the wild will be from the perspectives of both the builders and the users. Such trust is assessed at several points during

an iterative model building process. Nonetheless, many of the state-of-the-art machine learning-based models (*e.g.* neural networks) act as black-boxes once deployed. When such models are used for decision-making, the lack of explanations regarding why and how they have reached their decisions poses several concerns. In order to address this limitation, recent research efforts in the machine learning community have proposed different approaches to make the algorithms more amenable to *ex ante* and *ex post* inspection. For example, a number of studies have attempted to tackle the issue of discrimination within algorithms by introducing tools to both identify [6] and rectify [13, 6, 31] cases of unwanted bias. Recently, Ribeiro *et al.* [69] have proposed a model-agnostic method to derive explanations for the predictions of a given model.

An interesting ongoing initiative is the Open Algorithms (OPAL) project¹⁴, a multi-partner effort led by Orange, the MIT Media Lab, Data-Pop Alliance, Imperial College London, and the World Economic Forum, that aims to open -without exposing- data collected and stored by private companies by “sending the code to the data” rather than the other way around. The goal is to enable the design, implementation and monitoring of development policies and programs, accountability of government action, and citizen engagement while leveraging the availability of large scale human behavioral data. OPAL’s core will consist of an open platform allowing open algorithms to run on the servers of partner companies, behind their firewalls, to extract key development indicators and operational data of relevance for a wide range of potential users. Requests for approved, certified and pre-determined indicators by third parties –*e.g.* mobility matrices, poverty maps, population densities– will be sent to them via the platform; certified algorithms will run on the data in a multiple privacy-preserving manner, and results will be made available via an API. The platform will also be used to foster civic engagement of a broad range of social constituents –academic institutions, private sector companies, official institutions, non-governmental and civil society organizations. Overall, the OPAL initiative has three key objectives: (i) engage with data providers, users, and analysts at all the stages of algorithm development; (ii) contribute to building local capacities and help shaping the future technological, ethical and legal frameworks that will govern the collection, control and use of human behavioral data to foster social progress; and (iii) build data literacy among users and partners, conceptualized as “the ability to constructively engage in society through and about data”. Initiatives such as OPAL have the potential to enable more human-centric accountable and transparent data-driven decision-making and governance.

¹⁴ <http://datapopalliance.org/open-algorithms-a-new-paradigm-for-using-private-data-for-social-good/>

4.3 *Living labs to experiment data-driven policies*

The use of real-time human behavioral data to design and implement policies has been traditionally outside the scope of the way of working in policy making. However, the potential of this type of data will only be realized when policy makers are able to analyze the data, to study human behavior and to test policies in the real world. A possible way is to build living laboratories -communities of volunteers willing to try new ways of doing things in a natural setting- in order to test ideas and hypotheses in a real life setting. An example is the Mobile Territorial Lab (MTL), a living lab launched by Fondazione Bruno Kessler, Telecom Italia, the MIT Media Lab and Telefonica, that has been observing the lives of more than 100 families through multiple channels for more than three years [15]. Data from multiple sources, including smartphones, questionnaires, experience sampling probes, etc. has been collected and used to create a multi-layered view of the lives of the study participants. In particular, social interactions (*e.g.* call and SMS communications), mobility routines and spending patterns, etc. have been captured. One of the MTL goals is to devise new ways of sharing personal data by means of Personal Data Store (PDS) technologies, in order to promote greater civic engagement. An example of an application enabled by PDS technologies is the sharing of best practices among families with young children. How do other families spend their money? How much do they get out and socialize? Once the individual gives permission, MyDataStore [89], the PDS system used by MTL participants, allows such personal data to be collected, anonymized, and shared with other young families safely and automatically.

The MTL has been also used to investigate how to deal with the sensitivities of collecting and using deeply personal data in real-world situations. In particular, a MTL study investigated the perceived monetary value of mobile information and its association with behavioral characteristics and demographics; the results corroborate the arguments towards giving back to the people (users, citizens, according to the scenario) control on the data they constantly produce [77].

Along these lines, Data-Pop Alliance and the MIT Media Lab launched in May 2016 a novel initiative called “Laboratorio Urbano” in Bogotá, Colombia, in partnership with Bogotá’s city government and Chamber of Commerce. The main objective of the Bogotá Urban Laboratory is to contribute to the city’s urban vitality, with a focus on mobility and safety, through collaborative research projects and dialogues involving the public and private sectors, academic institutions, and citizens. Similar initiatives are being planned in other major cities of the global south, including Dakar, Senegal, with the goal of strengthening and connecting local ecosystems where data-driven innovations can take place and scale.

Figure 1 provides the readers with a visual representation of the factors playing a significant role in positive data-driven disruption.




Fig. 1 Requirements summary for positive data-driven disruption.

5 Conclusion

In this chapter we have provided an overview of both the opportunities and the risks of data-driven algorithmic decision-making for the public good. We are witnessing an unprecedented time in our history, where vast amounts of fine grained human behavioral data are available. The analysis of this data has the impact to help inform policies in public health, disaster management, safety, economic development and national statistics among others. In fact, the use of data is at the core of the 17 Sustainable Development Goals (SDGs) defined by United Nations, both in order to achieve the goals and to measure progress towards their achievement.

While this is an exciting time for researchers and practitioners in this new field of computational social sciences, we need to be aware of the risks associated with these new approaches to decision making, including violation of privacy, lack of transparency, information asymmetry, social exclusion and discrimination. We have proposed three human-centric requirements that we consider to be of paramount importance in order to enable positive disruption of data-driven policy-making: user-centric data ownership and management; algorithmic transparency and accountability; and living labs to experiment with data-driven policies in the wild. It will be only when we honor these requirements that we will be able to move from the feared tyranny of data and algorithms to a data-enabled model of democratic governance running against tyrants and autocrats, and for the people.

References

1. G.A. Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.
2. G.A. Akerlof and R.J. Shiller. *Animal spirits: How human psychology drives the economy, and why it matters for global capitalism*. Princeton University Press, 2009.
3. S. Barocas and A.D. Selbst. Big data’s disparate impact. *California Law Review*, 104:671–732, 2016.
4. L. Bengtsson, X. Lu, A. Thorson, R. Garfield, and J. Von Schreeb. Improved response to disasters and outbreaks by tracking population movements with mobile phone network data: a post-earthquake geospatial study in haiti. *PLoS Medicine*, 8(8), 2011.
5. Y. Benkler. *The wealth of networks*. Yale University Press, New Haven, 2006.
6. B. Berendt and S. Preibusch. Better decision support through exploratory discrimination-aware data mining: Foundations and empirical evidence. *Artificial Intelligence and Law*, 22(2):1572–8382, 2014.
7. V. D. Blondel, A. Decuyper, and G. Krings. A survey of results on mobile phone datasets analysis. *EPJ Data Science*, 4(10), 2015.
8. J. Blumenstock, G. Cadamuro, and R. On. Predicting poverty and wealth from mobile phone metadata. *Science*, 350(6264):1073–1076, 2015.
9. A. Bogomolov, B. Lepri, M. Ferron, F. Pianesi, and A. Pentland. Daily stress recognition from mobile phone data, weather conditions and individual traits. In *Proceedings of the 22nd ACM International Conference on Multimedia*, pages 477–486. 2014.
10. A. Bogomolov, B. Lepri, J. Staiano, E. Letouzé, N. Oliver, F. Pianesi, and A. Pentland. Moves on the street: Classifying crime hotspots using aggregated anonymized data on people dynamics. *Big Data*, 3(3):148–158, 2015.
11. A. Bogomolov, B. Lepri, J. Staiano, N. Oliver, F. Pianesi, and A. Pentland. Once upon a crime: Towards crime prediction from demographics and mobile data. In *Proceedings of the International Conference on Multimodal Interaction (ICMI)*, pages 427–434, 2014.
12. J. Burrell. How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 2016.
13. T. Calders and S. Verwer. Three naive bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2):277–292, 2010.
14. T. Calders and I. Zliobaite. Why unbiased computational processes can lead to discriminative decision procedures. In B. Custers, T. Calders, B. Schermer, and T. Zarsky, editors, *Discrimination and Privacy in the Information Society*, pages 43–57. 2013.
15. S. Centellegher, M. De Nadai, M. Caraviello, C. Leonardi, M. Vescovi, Y. Ramadian, N. Oliver, F. Pianesi, A. Pentland, F. Antonelli, and B. Lepri. The mobile territorial lab: A multilayered and dynamic view on parents daily lives. *EPJ Data Science*, 5(3), 2016.
16. S.P. Chainey, L. Tompson, and S. Uhlig. The utility of hotspot mapping for predicting spatial patterns of crime. *Security Journal*, 21:4–28, 2008.
17. A. Christin, A. Rosenblatt, and d. boyd. Courts and predictive algorithms. *Data & Civil Rights Primer*, 2015.
18. D.K. Citron and F. Pasquale. The scored society. *Washington Law Review*, 89(1):1–33, 2014.
19. K. Crawford and J. Schultz. Big data and due process: Toward a framework to redress predictive privacy harms. *Boston College Law Review*, 55(1):93–128, 2014.
20. M. De Choudhury, M. Gamon, S. Counts, , and E. Horvitz. Predicting depression via social media. In *Proceedings of the 7th International AAAI Conference on Weblogs and Social Media*, 2013.
21. P. De Filippi. The interplay between decentralization and privacy: The case of blockchain technologies. *Journal of Peer Production*, 7, 2015.

22. Y.-A. de Montjoye, C. Hidalgo, M. Verleysen, and V. Blondel. Unique in the crowd: The privacy bounds of human mobility. *Scientific Reports*, 3, 2013.
23. Y.-A. de Montjoye, L. Radaelli, V.K. Singh, and A. Pentland. Unique in the shopping mall: On the re-identifiability of credit card metadata. *Science*, 347(6221):536–539, 2015.
24. Y.-A. de Montjoye, E. Shmueli, S. Wang, and A. Pentland. Openpds: Protecting the privacy of metadata through safeanswers. *PloS One*, (10.1371), 2014.
25. R. de Oliveira, A. Karatzoglou, P. Concejero Cerezo, A. Armenta Lopez de Vicuña, and N. Oliver. Towards a psychographic user model from mobile phone usage. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*, pages 2191–2196. ACM, 2011.
26. S. Devarajan. Africa’s statistical tragedy. *Review of Income and Wealth*, 59(S1):S9–S15, 2013.
27. N. Diakopoulos. Algorithmic accountability: Journalistic investigation of computational power structures. *Digital Journalism*, 2015.
28. W. Easterly. *The Tyranny of Experts*. Basic Books, 2014.
29. J. Eck, S. Chainey, J. Cameron, and R. Wilson. Mapping crime: understanding hotspots. *National Institute of Justice: Washington DC*, 2005.
30. M. Faurholt-Jepsena, M. Frostb, M. Vinberga, E.M. Christensena, J.E. Bardram, and L.V. Kessinga. Smartphone data as objective measures of bipolar disorder symptoms. *Psychiatry Research*, 217:124–127, 2014.
31. M. Feldman, S.A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 259–268, 2015.
32. A.G. Ferguson. Crime mapping and the fourth amendment: Redrawing high-crime areas. *Hastings Law Journal*, 63:179–232, 2012.
33. G. Fields. Changes in poverty and inequality. *World Bank Research Observer*, 4:167–186, 1989.
34. S.T. Fiske. Stereotyping, prejudice, and discrimination. In D.T. Gilbert, S.T. Fiske, and G. Lindzey, editors, *Handbook of Social Psychology*, pages 357–411. Boston: McGraw-Hill, 1998.
35. E. Frias-Martinez, G. Williamson, and V. Frias-Martinez. An agent-based model of epidemic spread using human mobility and social network information. In *Social Computing (SocialCom), 2011 International Conference on*, pages 57–64. IEEE, 2011.
36. T. Gillespie. The relevance of algorithms. In T. Gillespie, P. Boczkowski, and K. Foot, editors, *Media technologies: Essays on communication, materiality, and society*, pages 167–193. MIT Press, 2014.
37. J. Ginsberg, M.H. Mohebbi, R.S. Patel, L. rammer, M.S. Smolinski, and L. Brilliant. Detecting influenza epidemics using search engine query data. *Nature*, 457:1012–1014, 2009.
38. Sara Hajian, Francesco Bonchi, and Carlos Castillo. Algorithmic bias: From discrimination discovery to fairness-aware data mining. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2125–2126. ACM, 2016.
39. N. Jean, M. Burke, M. Xie, W.M. Davis, D.B. Lobell, and S. Ermon. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794, 2016.
40. M. Jerven. *Poor numbers: How we are misled by african development statistics and what to do about it*. Cornell University Press, 2013.
41. G. King. Ensuring the data-rich future of the social sciences. *Science*, 2011.
42. M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15):5802–5805, 2013.

43. S. Kuznets. Economic growth and income inequality. *American Economic Review*, 45:1–28, 1955.
44. M. Latzer, K. Hollnbuchner, N. Just, and F. Saurwein. The economics of algorithmic selection on the internet. In J. Bauer and M. Latzer, editors, *Handbook on the Economics of the Internet*. Edward Elgar, Cheltenham, Northampton, 2015.
45. D. Lazer, A. Pentland, L. Adamic, S. Aral, A-L. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, T. Jebara, G. King, M. Macy, D. Roy, and M. Van Alstyne. Computational social science. *Science*, 323(5915):721–723, 2009.
46. B. Lepri, J. Staiano, E. Shmueli, F. Pianesi, and A. Pentland. The role of personality in shaping social networks and mediating behavioral change. *User Modeling and User-Adapted Interaction*, 26(2):143–175, 2016.
47. R. LiKamWa, Y. Liu, N.D. Lane, and L. Zhong. Moodscope: Building a mood sensor from smartphone usage patterns. In *Proceedings of the 11th Annual International Conference on Mobile Systems, Applications, and Service (MobiSys)*, pages 389–402. 2013.
48. H.Y. Liu, E. Skjetne, and M. Kobernus. Mobile phone tracking: In support of modelling traffic-related air pollution contribution to individual exposure and its implications for public health impact assessment. *Environmental Health*, 12, 2013.
49. T. Louail, M. Lenormand, O. G. Cantu Ros, M. Picornell, R. Herranz, E. Frias-Martinez, J. J. Ramasco, and M. Barthelemy. From mobile phone data to the spatial structure of cities. *Scientific Reports*, 4(5276), Jun 2014.
50. X. Lu, L. Bengtsson, and P. Holme. Predictability of population displacement after the 2010 haiti earthquake. *Proceedings of the National Academy of Sciences*, 109:11576–81, 2012.
51. B. Major and L.T. O’Brien. The social psychology of stigma. *Annual Review of Psychology*, 56:393–421, 2005.
52. A. Matic and N. Oliver. The untapped opportunity of mobile network data for mental health. In *Future of Pervasive Health Workshop*. ACM, 6 2016.
53. G.O. Mohler, M.B. Short, P.J. Brantingham, F.P. Schoenberg, and G.E. Tita. Self-exciting point process modeling of crime. *Journal of the American Statistical Association*, (106):100–108, 2011.
54. S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Technical report, Kent University, 2009.
55. F. Ofli, P. Meier, M. Imran, C. Castillo, D. Tuia, N. Rey, J. Briant, P. Millet, F. Reinhard, M. Parkan, and S. Joost. Combining human computing and machine learning to make sense of big (aerial) data for disaster response. *Big Data*, 4:47–59, 2016.
56. P. Ohm. Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review*, 57:1701–1777, 2010.
57. N. Oliver, A. Matic, and E. Frias-Martinez. Mobile network data for public health: Opportunities and challenges. *Frontiers in Public Health*, 3:189, 2015.
58. C. O’Neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown, 2016.
59. V. Osmani, A. Gruenerbl, G. Bahle, Lukowicz P. Haring, C., and Mayora O. Smartphones in mental health: Detecting depressive and manic episodes. *IEEE Pervasive Computing*, 14(3):10–13, 2015.
60. D. Pager and H. Shepherd. The sociology of discrimination: Racial discrimination in employment, housing, credit and consumer market. *Annual Review of Sociology*, 34:181–209, 2008.
61. F. Pasquale. *The Black Box Society: The secret algorithms that control money and information*. Harvard University Press, 2015.
62. D. Pastor-Escuredo, Y. Torres Fernandez, J.M. Bauer, A. Wadhwa, C. Castro-Correa, L. Romanoff, J.G. Lee, A. Rutherford, V. Frias-Martinez, N. Oliver, Frias-Martinez E., and M. Luengo-Oroz. Flooding through the lens of mobile phone activity. In *IEEE Global Humanitarian Technology Conference, GHTC’14*. IEEE, 2014.

63. A. Pentland. Society's nervous system: Building effective government, energy, and public health systems. *IEEE Computer*, 45(1):31–38, 2012.
64. W.L. Perry, B. McInnis, C.C. Price, S.C. Smith, and J.S. Hollywood. *Predictive policing: The role of crime forecasting in law enforcement operations*. Rand Corporation, 2013.
65. J. Podesta, P. Pritzker, E.J. Moniz, J. Holdren, and J. Zients. Big data: Seizing opportunities, preserving values. Technical report, Executive Office of the President, 2014.
66. E. Ramirez, J. Brill, M.K. Ohlhausen, and T. McSweeney. Big data: A tool for inclusion or exclusion? Technical report, Federal Trade Commission, January 2016.
67. J.H. Ratcliffe. A temporal constraint theory to explain opportunity-based spatial offending patterns. *Journal of Research in Crime and Delinquency*, 43(3):261–291, 2006.
68. M. Ravallion. *The economics of poverty: History, measurement, and policy*. Oxford University Press, 2016.
69. M.T. Ribeiro, S. Singh, and C. Guestrin. "why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, pages 1135–1144, 2016.
70. W. Samuelson and R. Zeckhauser. Status quo bias in decision making. *Journal of Risk and Uncertainty*, (1):7–59, 1988.
71. J. San Pedro, D. Proserpio, and N. Oliver. Mobiscore: Towards universal credit scoring from mobile phone data. In *Proceedings of the International Conference on User Modeling, Adaptation and Personalization (UMAP)*, pages 195–207, 2015.
72. M. B. Short, M. R. D'Orsogna, V. B. Pasour, G. E. Tita, P. J. Brantingham, A. L. Bertozzi, and L. B. Chayes. A statistical model of criminal behavior. *Mathematical Models and Methods in Applied Sciences*, 18(supp01):1249–1267, 2008.
73. V. K. Singh, B. Bozkaya, and A. Pentland. Money walks: Implicit mobility behavior and financial well-being. *PLOS ONE*, 10(8):e0136628, 2015.
74. V.K. Singh, L. Freeman, B. Lepri, and A. Pentland. Predicting spending behavior using socio-mobile features. In *Social Computing (SocialCom), 2013 International Conference on*, pages 174–179. IEEE, 2013.
75. C. Smith-Clarke, A. Mashhadi, and L. Capra. Poverty on the cheap: Estimating poverty maps using aggregated mobile communication networks. In *Proceedings of the 32nd ACM Conference on Human Factors in Computing Systems (CHI2014)*, 2014.
76. V. Soto, V. Frias-Martinez, J. Virseda, and E. Frias-Martinez. Prediction of socioeconomic levels using cell phone records. In *Proceedings of the International conference on UMAP*, pages 377–388, 2011.
77. J. Staiano, N. Oliver, B. Lepri, R. de Oliveira, M. Caraviello, and N. Sebe. Money walks: a human-centric study on the economics of personal mobile data. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 583–594. ACM, 2014.
78. J. Staiano, G. Zyskind, B. Lepri, N. Oliver, and A. Pentland. The rise of decentralized personal data markets. In D. Shrier and A. Pentland, editors, *Trust::Data: A New Framework for Identity and Data Sharing*. CreateSpace Independent Publishing Platform, 2016.
79. L. Sweeney. Discrimination in online ad delivery. Available at SSRN: <http://ssrn.com/abstract=2208240>, 2013.
80. N. Szabo. Formalizing and securing relationships on public networks. *First Monday*, 2(9), 1997.
81. L. Thomas. *Consumer credit models: Pricing, profit, and portfolios*. New York: Oxford University Press, 2009.

82. M. Tizzoni, P. Bajardi, A. Decuyper, G. Kon Kam King, C.M. Schneider, V. Blondel, Z. Smoreda, M.C. Gonzalez, and V. Colizza. On the use of human mobility proxies for modeling epidemics. *PLoS Computational Biology*, 10(7), 2014.
83. C. Tobler. Limits and potential of the concept of indirect discrimination. Technical report, European Network of Legal Experts in Anti-Discrimination, 2008.
84. J.L. Toole, N. Eagle, and J.B. Plotkin. Spatiotemporal correlations in criminal offense records. *ACM Transactions on Intelligent Systems and Technology*, 2(4):38:1–38:18, July 2011.
85. M. Traumueller, G. Quattrone, and L. Capra. Mining mobile phone data to investigate urban crime theories at scale. In *Proceedings of the International Conference on Social Informatics*, pages 396–411, 2014.
86. Z. Tufekci. Algorithmic harms beyond facebook and google: Emergent challenges of computational agency. *Colorado Technology Law Journal*, 13:203–218, 2015.
87. A. Tversky and D. Kahnemann. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.
88. A. Venerandi, G. Quattrone, L. Capra, D. Quercia, and D. Saez-Trumper. Measuring urban deprivation from user generated content. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW2015)*, 2015.
89. M. Vescovi, C. Perentis, C. Leonardi, B. Lepri, and C. Moiso. My data store: Toward user awareness and control on personal data. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, pages 179–182, 2014.
90. H. Wang, Z. Li, D. Kifer, and C. Graif. Crime rate inference with big data. In *Proceedings of International conference on KDD*, 2016.
91. T. Wang, C. Rudin, D. Wagner, and R. Sevieri. Learning to detect patterns of crime. In *Machine Learning and Knowledge Discovery in Databases*, pages 515–530. Springer, 2013.
92. R. Want, T. Pering, G. Danneels, M. Kumar, M. Sundar, and J. Light. The personal server: Changing the way we think about ubiquitous computing. In *Proceedings of 4th International Conference on Ubiquitous Computing*, pages 194–209, 2002.
93. D. Weisburd. Place-based policing. *Ideas in American Policing*, 9:1–16, 2008.
94. A. Wesolowski, N. Eagle, A. Tatem, D. Smith, R. Noor, and C. Buckee. Quantifying the impact of human mobility on malaria. *Science*, 338(6104):267–270, 2012.
95. A. Wesolowski, G. Stresman, N. Eagle, J. Stevenson, C. Owaga, E. Marube, T. Bousema, C. Drakeley, J. Cox, and C.O. Buckee. Quantifying travel behavior for infectious disease research: A comparison of data from surveys and mobile phones. *Scientific Reports*, 4, 2014.
96. M. Willson. Algorithms (and the) everyday. *Information, Communication & Society*, 2016.
97. R. Wilson, E. Erbach-Schoenengerg, M. Albert, D. Power, Tudge S., and Gonzalez M. et al. Rapid and Near Real-time Assessments of Population Displacement Using Mobile Phone Data Following Disasters: The 2015 Nepal Earthquake. *PLOS Current Disasters*, February 2016.
98. H. Zang and J. Bolot. Anonymization of location data does not work: A large-scale measurement study. In *Proceedings of 17th ACM Annual International Conference on Mobile Computing and Networking*, pages 145–156, 2011.
99. T. Zarsky. The trouble with algorithmic decisions: An analytic road map to examine efficiency and fairness in automated and opaque decision making. *Science, Technology, and Human Values*, 41(1):118–132, 2016.
100. T.Z. Zarsky. Automated prediction: Perception, law and policy. *Communications of the ACM*, 4:167–186, 1989.
101. G. Zyskind, O. Nathan, and A. Pentland. Decentralizing privacy: Using blockchain to protect personal data. In *Proceedings of IEEE Symposium on Security and Privacy Workshops*, pages 180–184. 2014.