University of Groningen

Caveat Emptor

Latif, Siddique; Qayyum, Adnan; Usama, Muhammad; Qadir, Junaid; Zwitter, Andrej; Shahzad, Muhammad

Published in:
 IEEE technology and society magazine

DOI:
[10.1109/MTS.2019.2930273](https://doi.org/10.1109/MTS.2019.2930273)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
 Publisher's PDF, also known as Version of record

Publication date:
 2019

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Latif, S., Qayyum, A., Usama, M., Qadir, J., Zwitter, A., & Shahzad, M. (2019). Caveat Emptor: The Risks of Using Big Data for Human Development. *IEEE technology and society magazine*, 38(3), 82-90.
<https://doi.org/10.1109/MTS.2019.2930273>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

ISTOCK/AGARTOFILM

Siddique Latif, Adnan Qayyum, Muhammad Usama, Junaid Qadir,
Andrej Zwitter, and Muhammad Shahzad

Caveat Emptor

The Risks of Using Big Data for Human Development

Big Data” has the potential to facilitate sustainable development in many sectors of life such as education, health, agriculture, and in combating humanitarian crises and violent conflicts. However, lurking beneath the immense promises of Big Data are some significant risks such as 1) the potential use of Big Data for unethical ends; 2) its ability to mislead through reliance on unrepresentative and biased data; and 3) the various privacy and security challenges associated with data (including the danger of an adversary tampering with the data to harm people). These risks can have severe consequences and a better understanding of these risks is the first step towards their mitigation

Digital Object Identifier 10.1109/MTS.2019.2930273
Date of publication: 30 August 2019

of these risks. In this article, we highlight the potential dangers associated with using Big Data, particularly for human development.

Data Deluge

Over the last decades, widespread adoption of digital applications has moved all aspects of human lives into the digital sphere. The commoditization of the data collection process due to increased digitization has resulted in a “data deluge” that continues to intensify with a number of Internet companies dealing with petabytes of data on a daily basis. The term “Big Data” has been coined to refer to our emerging ability to collect, process, and analyze the massive amount of data being generated from multiple sources in order to obtain previously inaccessible insights. Big Data can equip policy and decision makers with evidence-based actionable insights that can help in enhancing social systems, tracking development progress, and in developing a nuanced understanding of the effects of policies without being swayed by intuition, ideologies, or emotions.

In particular, the recent advances in Machine Learning (ML) and Artificial Intelligence (AI) techniques have revolutionized intelligent data analytics resulting in enhanced interest in using Big Data for (sustainable human and social) development (BD4D). Many organizations and government institutions are exploring such solutions in diverse fields such as healthcare, education, intelligence, fraud, and crime prevention (5).

Although Big Data technology offers great promise (3), it is worth remembering that Big Data is not a silver bullet and we may ignore the hard-earned statistical lessons on measurement bias, data quality, and inference variation that have been earned through a hard toil and sometimes, bitter experience only at our own peril. While most writing on data is enthusiastic, new work has started emerging that has begun to show how Big Data can mislead and be used detrimentally (7), (20), (36). More than 50 years of research into artificial intelligence and statistical learning has shown that there is no free lunch — i.e., there are no universally applicable solutions (50) and that there are always trade-offs involved (47).

The use of BD4D is transforming society in various domains from education to prevention, diagnosis, and treatment of illness (16). It has improved the efficiency and effectiveness of disaster management systems by utilizing real-time community information (39). In these applications, data is emerging as a new economic resource and used by companies, governments, and even individuals to optimize everything (31). The ultimate goal of the development sector seems to be what is referred to as forecast-based financing — a

prediction-based and automatized funding and logistical system based on big data analytics and smart contracts that automatizes everything from funding to action in the field (53).

Despite the great excitement around the Big Data trend, Big Data is also pegged with criticism. The dark side of Big Data is that data invariably contains some biases and there is a fear that Big Data could erode privacy and threaten freedom when deployed for human and social development. These issues become particularly serious, when Big Data is deployed for human subjects in the case of BD4D. For example, consider the case where Big Data predictions about individuals are used to punish people on their propensities rather than on their actions. This denies human free will and their capability to improve over time and effectively reinforces existing stereotypes (24).

Various studies have articulated common issues that concern Big Data analytics interdisciplinary research (13), (19), (33), (44) but a critical analysis of BD4D is missing in the literature. Therefore, in this paper, we discuss the caveats of Big Data analytics when it is used for human development purposes. The purpose of this article is to present a critical synthesis of the diverse literature and put together various underlying issues that affect BD4D significantly. In addition to critically evaluating BD4D risks, this paper also explores potential remedies that can help mitigate the problems of BD4D.

General Issues in Using Data for Decisions

Today, big corporations are investing their resources in utilizing Big Data technology to uncover important correlations, customer preferences, market trends, and other hidden patterns that can help boost their production efficiency by providing data-driven products and services. Despite these great opportunities, Big Data has come with a raft of potential challenges and pitfalls that need to be considered. Some of them are discussed next.

Big Data is Reductionist

Big Data provides a granular view of extremely complex and historically unprecedented problems by using particular vantage points from the available data. It only focuses on specific types of questions, and seeks their answers in a reductionist way by seemingly ignoring concrete contextual realities about human societies and places. It is also crude and reductionist in its techniques and algorithms by sacrificing complexity, specificity, and deep contextual knowledge (26). In this way, analyses are reduced to simple outputs that critically ignore the underlying complexity of the actual processes and entities being modeled.

Big Data is not Neutral nor is it Objective

One of the compelling aspects of Big Data is the assumption that it is eliminating human subjectivity and bias. Although the idea of “baggage-free” learning is seductive, data alone is not enough since there is no learning without knowledge as a practical consequence of the “no free lunch” theorem. Big Data analytics is what mathematicians call an ill-posed problem, thus no unique model exists¹; the only way to resolve some ill-posed problems is to make additional assumptions. These assumptions are subjective and can be thought of as being “opinions embedded in Math” (36). Therefore, while data is often considered neutral, data collection techniques and analysis methods are designed by humans based on their knowledge, experience, and beliefs.

An uncritical acceptance and careless application of Big Data methods can adversely affect human welfare and development.

The German philosopher Heidegger wrote,

“Everywhere we remain unfree and chained to technology, whether we passionately affirm or deny it. But we are delivered over to it in the worst possible way when we regard it as something neutral.”

Big Data scientists make computations using data and present results about people in terms of numbers or mathematical relations based on layers upon layers of assumptions. Data scientists use this data and try to organize it into objects, relations, and events, and in doing so invariably analyze the data within the backdrop of their subjectivity and partiality. Moreover, it is challenging to ensure fairness in ML algorithms due to significant statistical limitations and it is possible that enforcing fairness can harm the very groups that these algorithms were supposed to protect (10).

Opaque Black-Box Models

Many of the modern mathematical models used for Big Data analysis are opaque — and unregulated (36). These models are often used inappropriately and can scale up biases thereby producing misleading results.

Such models have even been referred to as “weapons of maths destruction” (WMD) due to their inscrutable black box nature. These WMDs can cause havoc/harm and tend to punish the poor due to vicious feedback loops (36). The reliance on opaque algorithmic black boxes for decision making is even more problematic when we consider how easy it is for adversaries to attack ML models. New research attacks have emerged in which adversaries can trick ML models to make decisions as they desire — something that has huge consequences for a society in which human decisions are automated and made by machines.

Big Data’s Big Bias Problem

Most Big Data datasets contain hidden biases both in collection and analysis stage. These biases in data can create an illusion of reality. These biases are hard to undo and their elimination have unintended consequences on the results. Four major biases of Big Data are described as follows.

Sampling Bias: When the samples are partial, selective, and not random, the patterns of omitted information may influence the structures discovered in the data. Such samples will be unable to accurately predict the outcomes and will have reduced validity (38). Such bias is often found when we focus on sentiment analysis on social media, safety monitoring using social media, or use population statistics and tourism statistics.

Activity Bias: This is another bias that is usually encountered in web data (29). This arises from the time-based correlation of users’ activities across diverse websites, because most of the users visit a website for a while and never return during the measurement period. This can be explained using the example of the model of Ginsberg *et al.* (18) for predicting flu cases. Using 50 million Google search terms, the model suffered from activity bias where it started over-predicting the flu in the U.S. (9), (28).

Information Bias: This refers to the delusion that more information always results in better decisions. Data scientists give too much significance on the volume of data and ignore other alternative opinions and solutions. The irony of this belief is that fewer data can give better decisions in some situations. In various cases, a simple rule of thumb can work better than complex analysis. There are situations where ignorance from very large data and required calculation can provide more accurate and faster results because large biased data just magnify the errors. Having more information is not always more desirable.

Inductive Bias (Assuming the Future Will be Like the Past): In Big Data analysis, there is an implicit belief that the future can always be extrapolated based on historical data. Apart from this approach being philosophically

¹A famous statistical quote by George Box is, “All models are wrong, some are useful.”

debatable (cf., the “Problem of Induction”), such an approach can empirically backfire in many cases and result in invalid, misleading, or unhelpful analyses.

Fooled by the “Big Noise”

In (43), the author Nate Silver notes that even though data is increasing globally by 2.5 quintillion bytes each day, useful information is of course not increasing commensurately. This means that most of the data is just noise and that noise is increasing faster than useful information. A related problem is that of the curse of dimensionality, where the more dimensions one works with, the less effective standard computational and statistical techniques become, which has serious repercussions when we are dealing with Big Data. As more data is generated and collected, the complexity, deviations, variance (or noise) and the number of potential false findings grow exponentially compared to the information (or signal). This spurious rise in data can mislead us to fake statistical relationships.

Another point of concern is that a plethora of hypotheses are tested using a single data in Big Data analytics, which opens the door for *spurious correlations*. When a large number of hypotheses are tested, it becomes highly likely that some false results will become statistically significant, which can mislead Big Data practitioners if they are not careful. The proper way to formulate a hypothesis is before the experiment not after it. It is necessary to understand that the statistical significance under an inaccurate method is totally specious — i.e., significance tests do not protect against “data dredging” (45).

Apart from the problem of spurious correlations, Big Data analysts can also become guilty of *cherry picking*. In this phenomenon, scientists tend to focus on finding support for their hypotheses and beliefs while overlooking other evidence. As a result, they only present positive results of their experiments that endorse their hypothesis or argument instead of reporting all of the findings (34). In this way, Big Data provide a relatively small advantage to improve public strategies, policies, and operations with minimum public benefit (8).

Are Big Data Predictions Generalizable?

In Big Data analytics, datasets are mostly gathered from observational sources instead of a strict statistical experiment. This poses the question of the generalizability of the insights learned from this data, notwithstanding the large size of the dataset. Even though we know that statistically speaking the discrepancy in an estimate decreases with an increase in the sample size as Bernoulli proved, but the messiness of real-world data in terms of incompleteness,

imbalance, hidden bias, rareness, large variances, outliers, and non-independent, identically distributed (IID) nature, means that simply getting more data is not sufficient. Disregard for the messiness of the real-world data can mislead us into problems such as multiple testing, regression to the mean, selection bias, and over-interpretation of causal associations (36). Also we must note that certain things are not precisely predictable (e.g., chaotic processes, complex systems, the so-called black swans) no matter how much data is available.

Big Data provides opportunities for unscrupulous people, groups, government, or organizations to use users’ personal data against them.

In BD4D applications, large administrative databases may have obscure/indeterminate data quality, limited information about confounding variables, and suboptimal documentation of the outcome measures; careful attention should be given to the generalizability of the learned insights.

Data in the Field of Development

Ethical Use of Data and Privacy Issues

Big Data opens frightening opportunities for unscrupulous people, group, government, or organizations to use users’ personal data against them for evil purposes as various facets of human lives are being digitized (7). Privacy depends on the nature and characteristics of data, the context in which it was created, and the norms and expectations of people (7). Therefore, it is necessary to situate and contextualize data in such a way that minimizes privacy breaches without affecting its usefulness.

Big Data can bring more transparency both for individuals and organizations by introducing user anonymization but there are numerous examples when data is thought to be anonymous but its combination with other variables resulted in unexpected re-identification (2), (27), (42). For example, it has been proved that four spatiotemporal points are sufficient to uniquely identify 95% of the individuals in a dataset where the location of an individual is specified hourly with a spatial resolution given by the carrier’s antennas (11). In another similar study, Latanya Sweeney (46) argued that 87% of all

Americans could be uniquely recognized using only three bits of information: zip code, birth date, and sex. The datasets being released today are anonymized by applying *ad hoc* de-identification methods. Therefore, the possibility of re-identification depends heavily on the advancement of re-identification methods and availability of the auxiliary datasets to an adversary. Thus, the chances of privacy breaches in the future are essentially uncertain.

The all-important consideration to be accounted for is how the collection of data may affect a person's well-being and dignity.

The all-important consideration to be accounted for is how the collection of data may affect a person's well-being and dignity, and to ensure that basic human rights and human freedom are not impinged upon. We should work actively to minimize potential risks related to the release of private confidential information, and the malicious use of data (potentially by unauthorized agents).

Quality Of Statistics — Numbers are Soft, and Incentives Matter

Data is not simply information that is harvested from an objective context. It should be an institutionally, culturally, and socially shaped product. Collecting good quality data for development is in practice very costly as we need resources and manpower for data collection, maintenance, and monitoring (22). For instance, the aggregate statistic for Gross Domestic Product (GDP) in many sub-Saharan African countries is often measured approximately since most African countries are simply unable to collect all the information needed to calculate GDP accurately and changes are approximately inferred from rainfall figures or population growth (21). The use of this data can lead to distorted or misleading policy decisions by development agencies and governments and can contribute to the underestimation of GDP and bewildering fallouts such as the following:

- 1) In November 2010, Ghana's Statistics Service announced an estimated GDP that was off from its true value, resulting in an upward adjustment of a whopping 60%, which was enough to change Ghana's status from a low-income country to a lower middle-income country (23);

- 2) Similar adjustments were made in Nigeria in April 2014 where the rise was almost 90%, which caused the total GDP of sub-Saharan Africa to rise by 20% (21).

But where did these overnight growth spurts come from? And what to say about the fate of the analysis and policies of the various policymakers and development professionals that were based on the previous miscalculated data — and the cost of this misplaced analysis? The official statistics are often missing, incomplete, dated, or unreliable. Given such concerns, researchers have argued that the numbers cannot be taken at face value for development data from a place like Africa, resulting in what has been called "Africa's statistical tragedy" (21).

Pitfalls of Self-Monitoring of States

In the modern world, almost every government tightly controls what kind of information is disseminated about the state. An important aspect of the self-monitoring of states is that it often becomes a farce because people cannot keep their own score. The official statistics measured by the government must be accepted with a pinch of salt since governments are spinning these numbers in ways that project the country's progress positively since the legitimacy of the state depends on the popular understanding of the country's progress. The anthropologist James Scott in his book "Seeing Like a State" (41) describes the ways in which governments, in their fetish for quantification and data, end up making people's lives miserable rather than better. As development projects initiated by the states are aimed at improving human lives, an appropriate measure for the development must be how much effort is expended on minimizing the risks associated with data collection and analysis.

Human Systems Have Complex Loops with Predictions Being Self-Fulfilling

Human social systems are complex adaptive systems with multi-loop nonlinear feedback in which actions performed for some purpose can lead to quite different unexpected and unintended consequences. We will use two examples to illustrate the 1) unintended consequences and 2) self-fulfilling nature of interventions in complex social systems.

For our first example, consider the Cobra effect as an example of an incident in which an intervention crafted to ameliorate a problem actually aggravated it by producing some unintended consequences. During the British colonial rule of India, the government devised a bounty system for combating the rise of venomous cobras. The system worked successfully initially and lots of snakes were killed for the reward but entrepreneurs

soon figured out that they could make money by farming cobras and killing more of them. The government, on learning about this, scrapped the system but ended up with a situation in which there were more cobras after the intervention than before.

For our second example, consider the Paper Town effect. In the 20th century, a famous map of New York was created by cartographers Lindberg and Alpers, who cleverly embedded a fake city Agloe into their map. Agloe, NY, was not a real town. It was a paper town — a booby trap to catch plagiarizers. A few years after Lindberg and Alpers set their map trap, the fake town appeared on another map by the cartographer Rand McNally, prompting the two mapmakers to sue for copyright infringement. Eventually it was discovered that a real town called Agloe had in fact emerged in New York² since users of the Linberg Alpers map thought that the Agloe that once existed must have gone missing and rebuilt it, and that Rand McNally may not have after all have plagiarized Lindberg and Alpers.

Missing Data Problem

Missing data is a big problem for development statistics. It has been reported by Jerven that around half of the 82 low-income countries have had one or partial poverty surveys within the past decade (21). Kaiser Fung, in his book *Numbersense* (15), discourages the assumption that we have everything and says: “N = All is often an assumption rather than a fact about the data.” Due to missing data problem, any ranking of such countries based on GDP only will be misleading because of the uneven use of methods and access to data. Handling missing data is very important for data mining processes as missing observations can significantly affect the performance of the model (32). Therefore, analysts should handle the missing patterns by employing appropriate methods to cope with it and to avoid the “streetlight effect.” The streetlight effect is the trend adopted by researchers to study what is easy to study. The streetlight effect is a major issue that keeps Big Data findings from being realistically useful for human development — especially when findings are yielded using user-generated and easily available data.

Some Remedies

Interpretable AI and Big Data Analysis

With the wide adoption of deep learning and ensemble methods, modern AI systems have become complex and opaque and increasingly operate as black boxes. Although these black-box models are producing

outstanding results, it is very difficult to trust their predictions due to their opaque and hard-to-interpret nature. This has been dubbed by some as the AI interpretability problem. We can define interpretability as the ability to describe the internal processes of a system (i.e., complex AI and ML techniques) in such a way that they are understandable to humans (17).

Understanding the internal workings of modern AI systems can help ensure fairness, bias identification, robustness, and generalization of Big-Data-based AI models.

Interpretable AI can help ensure algorithmic fairness, bias identification, robustness, and generalization of Big-Data-based AI models. Interpretation of AI-based decision-making algorithms is also necessary to ensure smooth deployment of real-world intelligent systems. But the development of interpretable AI requires that the following questions be answered:

- 1) How to ensure accountability of a model?
- 2) How to ensure the transparency of the model output?
- 3) How to ensure the fairness of the model predictions?

Since BD4D is directly related to human development, the Big-Data-based AI model used in these cases (i.e., for healthcare systems, judicial systems, etc.) must ensure high accuracy and interpretability. A possible remedy is to insist on using interpretable ML for high-stakes BD4D decisions and to utilize explanation methods for justifying the decisions, where explanation means the provision of visual or textual evidence of a certain features related to an AI model’s decision. As an example of work in this space, Bach *et al.* (6) proposed the layer-wise relevance propagation (LRP) method which provides a visual contribution for each input feature in decision making. In another work, Ribeiro *et al.* (40) proposed two explanation methods namely locally interpretable model-agnostics (LIME) and submodular pick locally interpretable model-agnostics (SP-LIME). Interpretable AI for BD4D is still an open research avenue to explore for Big Data and the AI community.

Better Generalization

In Big Data research, researchers make inferences by training models on a larger subset of data with the goal

²https://en.wikipedia.org/wiki/Agloe,_New_York

to fit the learned hypothesis on unseen data. In Big Data research, generalization is the procedure of spanning the characteristics of a group or class to the entire group or population. This enables the inference of attributes of an entire population without getting to know every single element in that population individually. The problem comes along when we wrongly generalize; or more precisely when we overdo it. The generalization fallacy occurs when statistical inferences about a particular population are asserted to a group of people for which the original population is not a representative sample. In other words, models overfit when they learn not only the signal from the training data but also the noise that impedes the model's capability to predict on unseen data. In order to avoid excessive generalization error, researchers should also check the scope of the results instead of extending scientific findings to the

"I understand that my work may have enormous effects on society and the economy, many of them beyond my comprehension."

whole population. Regularization is helpful to avoid overfitting by reducing the number of parameters to fit model in high-dimensional data (35). It also prevents model parameters to change easily, which helps in keeping the focus of the model on the persistent structure. Apart from regularization, other techniques such as cross-validation, early stopping, weight sharing, weight restriction, sparsity constraints, etc., can also be used for reducing the generalization error based on the algorithm being used.

Avoiding Bias

Big Data tends to have high dimensionality and may be conflicting, subjective, redundant, and biased. The awareness of potential biases can improve the quality of decisions at the level of individuals, organizations, and communities (25). In a study of 1000 major business investments conducted by McKinsey, it was found that when organizations worked to minimize the biases in their decision-making, they achieved up to 7% higher returns (30). The biases associated with multiple comparisons can be deliberately avoided using techniques such as the Bonferroni correction (48), the Sidak correction, and the HolmBonferroni correction (1). Another source of bias is called data snooping or data dredging,

which occurs when a portion of data is used more than once for model selection or inference. In technical evaluations of results, it is conceivable to repeat experiments using the same dataset to get satisfactory results (49). Data dredging can be avoided by conducting randomized out-of-sample experiments during hypotheses building. For example, an analyst gathers a dataset and arbitrarily segments it into two subsets, A and B. Initially, only one subset — say, subset A — is analyzed for constructing hypotheses. Once a hypothesis is formulated, it should then be tested on subset B. If subset B also supports such a hypothesis, then it might be trusted as valid. Similarly, we should use such models that can consider the degree of data snooping for obtaining genuinely good results.

Finding Causality Rather than Correlations

In most data analysis performed in the Big Data era, the focus is on determining correlations rather than on understanding causality (37). For BD4D problems, we're more interested in determining causes rather than correlates and therefore we must place a premium on performing causal BD4D analysis since causally driven analysis can improve BD4D decisions. Discovering causal relations is difficult and involves substantial effort, and requires going beyond mere statistical analysis as pointed out by Freedman (14) who has highlighted that for data analytics to be practically useful, it should be problem-driven or theory driven, not simply data-driven. As Freedman says, using Big Data for development requires "the expenditure of shoe leather" to situate the work in the proper context. The focus on correlation has arisen because of the lack of a suitable mathematical framework for studying the slippery problem of causality until the recent fundamental progress made by Pearl (37), whose work has now provided a suitable notation and algebra for performing a causal analysis.

Stress High-Quality Data Analytics Rather Than Big Data Analytics

A better and thoughtful understanding of risks or pitfalls of Big Data is crucial to decrease its associated potential harms to individuals and society. There needs to be a stress on utilizing Big Data along with data collected through traditional sources to provide a deeper, clearer understanding of problems, instead of being fixated on only generating and analyzing large volumes of data. Although it is generally preferred to have more data, it is not always desirable, especially in the cases where data is biased. Another disadvantage of large datasets is cost in terms of processing, storage, and maintenance. However, some simple methods like sampling and/or resampling enable us to extract the most relevant data from a larger chunk of data. Another very important aspect is

to collect the desired data to properly design an experiment rather than collecting all possible information. Specifically, in the field of human development and humanitarian action, there is no way around corroborating findings based on Big Data with intelligence gathered at the field level. This requires that international organizations and development actors actively increase their capacity for data collection and analysis also referred to as “Humanitarian Intelligence” (52).

User-Friendly and Responsible Data Analytics

Big Data algorithms are not as trustworthy as typically assumed since they draw upon data collected from a prejudiced and biased world (4). Cathy O’Neil (36) describes how algorithms often perpetuate or worsen inequality and injustice and suggests that there should be laws and industry standards to ensure transparency for Big Data gathering and utilization. In particular, false, outdated, and taken out of context information may cause harm to an individual’s autonomy and the use of such information should, therefore, be restricted. As a remedy to this issue, “the right to be forgotten” enables data subjects to reassert control over their personal information. There might be fair audits of algorithms but first, the awareness of this issue to programmers is required as they share a disproportionate amount of responsibility in the design of Big-Data-based algorithms. Since AI and ML algorithms are embedded in many crucial social systems — ranging from crime fighting to job portals to hospitals — it is recommended that social-systems analysis should also include the possible effects of AI on their performance throughout the various stages of the design cycle (i.e., conception, design, implementation, and deployment) in social institutions.

BD4D practitioners should follow the lead of the following five principles of data for humanity laid down in (51):

- 1) Do no harm;
- 2) Use data to help create peaceful coexistence;
- 3) Use data to help vulnerable people and people in need;
- 4) Use data to preserve and improve the natural environment; and
- 5) Use data to help create a fair world without discrimination.

BD4D practitioners will also do well to adhere to the following oaths, which were developed by Herman and Wilmott as the “Modelers’ Hippocratic Oath” (12) in the light of the global financial crisis:

- 1) Though I will use models boldly to estimate value, I will not be overly impressed by mathematics;
- 2) I will never sacrifice reality for elegance without explaining why I have done so;

- 3) Nor will I give the people who use my model false comfort about its accuracy. Instead, I will make explicit its assumptions and oversights; and
- 4) I understand that my work may have enormous effects on society and the economy, many of them beyond my comprehension.

Cautious Perspective on the Use of Big Data

In this paper, we have provided a cautious perspective on the use of Big Data for human development. While we believe that Big Data has great potential for facilitating human development, our aim is to caution against an uncritical acceptance and careless application of Big Data methods in matters directly affecting human welfare and development. We need to guard against a naïve overreliance on data to avoid the many pitfalls of data worship. We argue that Big Data technology is a tool, and like all tools it should be considered as a handmaiden rather than as a headmaster. In particular, we argue that Big Data analytics cannot substitute for good research design and subject-matter knowledge. Various potential remedies to address the pitfalls of using Big Data for Human and Social Development (BD4D) have also been highlighted. To conclude, we will like to emphasize that our paper should certainly not be construed as a techno-phobic manifesto. We believe strongly in the promise of BD4D but when pursued with due attention to the mitigation of the many associated pitfalls.

Author Information

Siddique Latif is with the University of Southern Queensland, Australia, and with Information Technology University (ITU)-Punjab, Pakistan.

Adnan Qayyum, Muhammad Usama, and Junaid Qadir are with Information Technology University (ITU)-Punjab, Pakistan.

Andrej Zwitter is with University of Groningen, Netherlands.

Muhammad Shahzad is with National University of Sciences and Technology (NUST), Pakistan.

References

- (1) H. Abdi, “Holm’s sequential Bonferroni procedure,” *Encyclopedia of Res. Design*, vol. 1, no. 8, pp. 1-8, 2010.
- (2) A. Acquisti, R. Gross, and F. Stutzman, “Privacy in the age of augmented reality,” *Proc. National Acad. Sciences*, 2011.
- (3) A. Ali, J. Qadir, R. Rasool, A. Sathaseelan, A. Zwitter, and J. Crowcroft, “Big data for development: applications and techniques,” *Big Data Analytics*, vol. 1, no. 1, p. 2, 2016.
- (4) P. Andras, L. Esterle, M. Guckert, T.A. Han, P.R. Lewis, K. Mila-novic, T. Payne, C. Perret, J. Pitt, S.T. Powers *et al.*, “Trusting intelligent machines: Deepening trust within socio-technical systems,” *IEEE Technology and Society Mag.*, vol. 37, no. 4, pp. 76-83, 2018.
- (5) P.M. Asaro, “AI ethics in predictive policing: From models of threat to an ethics of care,” *IEEE Technology and Society Mag.*, vol. 38, no. 2, pp. 40-53, June 2019.
- (6) S. Bach, A. Binder, G. Montavon, F. Klauschen, K.R. Muller, and W. Samek, “On pixel-wise explanations for non-linear classifier

- decisions by layer-wise relevance propagation," *PLoS one*, vol. 10, no. 7, p. e0130140, 2015.
- (7) D. Boyd and K. Crawford, "Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon," *Information, communication & Society*, vol. 15, no. 5, pp. 662–679, 2012.
- (8) D. Brooks, "The philosophy of data," *New York Times*, vol. 4, 2013.
- (9) D. Butler, "When Google got flu wrong," *Nature*, vol. 494, no. 7436, p. 155, 2013.
- (10) S. Corbett-Davies and S. Goel, "The measure and mismeasure of fairness: A critical review of fair machine learning," *arXiv preprint arXiv:1808.00023*, 2018.
- (11) Y.-A. De Montjoye, C.A. Hidalgo, M. Verleysen, and V.D. Blondel, "Unique in the crowd: The privacy bounds of human mobility," *Scientific Rep.*, vol. 3, p. 1376, 2013.
- (12) E. Derman and P. Wilmott, "The financial modelers' manifesto," 2009; https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1324878.
- (13) H. Ekbil, M. Mattioli, I. Kouper, G. Arave, A. Ghazinejad, T. Bowman, V.R. Suri, A. Tsou, S. Weingart, and C.R. Sugimoto, "Big data, bigger dilemmas: A critical review," *J. Association for Information Science and Technology*, vol. 66, no. 8, pp. 523–1545, 2015.
- (14) D.A. Freedman, "Statistical models and shoe leather," *Sociological Methodology*, pp. 291–313, 1991.
- (15) K. Fung, *Numbersense: How to Use Big Data to Your Advantage*. McGraw-Hill, 2013.
- (16) N. Giatrakos, N. Katzouris, A. Deligiannakis, A. Artikis, M. Garofalakis, G. Paliouras, H. Arndt, R. Grasso, R. Klinkenberg, M. Ponce De Leon, G.G. Tartaglia, A. Valencia, and D. Zisis, "Interactive extreme: Scale analytics towards battling cancer," *IEEE Technology and Society Mag.*, vol. 38, no. 2, pp. 54–61, June 2019.
- (17) L.H. Gilpin, D. Bau, B.Z. Yuan, A. Bajwa, M. Specter, and L. Kagal, "Explaining explanations: An approach to evaluating interpretability of machine learning," *arXiv preprint arXiv:1806.00069*, 2018.
- (18) J. Ginsberg, M. H. Mohebbi, R.S. Patel, L. Brammer, M.S. Smolinski, and L. Brilliant, "Detecting influenza epidemics using search engine query data," *Nature*, vol. 457, no. 7232, p. 1012, 2009.
- (19) I. Abaker, T. Hashem, I. Yaqoob, N.B. Anuar, S. Mokhtar, A. Gani, and S.U. Khan, "The rise of big data on cloud computing: Review and open research issues," *Information Systems*, vol. 47, pp. 98–115, 2015.
- (20) D. Helbing, B.S. Frey, G. Gigerenzer, E. Hafen, M. Hagner, Y. Hofstetter, J. van den Hoven, R.V. Zicari, and A. Zwitter, "Will democracy survive big data and artificial intelligence," *Scientific American*, vol. 25, 2017.
- (21) M. Jerven, *Poor Numbers: How We are Misled by African Development Statistics and What to Do About It*. Cornell Univ. Press, 2013.
- (22) M. Jerven, "How much will a data revolution in development cost?," in *Forum for Development Studies*, vol. 44. Taylor & Francis, 2017, pages 31–50.
- (23) M. Jerven and M.E. Duncan, "Revising GDP estimates in sub-Saharan Africa: Lessons from Ghana," *African Statistical J.*, vol. 15, pp. 13–24, 2012.
- (24) S.J. Walker, "Big data: A revolution that will transform how we live, work, and think," *Int. J. Advertising*, vol. 33, no. 1, pp. 181–183, 2014.
- (25) D. Kahneman, D. Lovallo, and O. Sibony, "Before you make that big decision," *Harvard Bus. Rev.*, vol. 89, no. 6, pp. 50–60, 2011.
- (26) R. Kitchin, "Big data, new epistemologies and paradigm shifts," *Big Data & Society*, vol. 1, no. 1, p. 2053951714528481, 2014.
- (27) I.M. Kloumann and J.M. Kleinberg, "Community membership identification from small seed sets," in *Proc. 20th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*. ACM, 2014, pp. 1366–1375.
- (28) D. Lazer, R. Kennedy, G. King, and A. Vespignani, "The parable of Google Flu: traps in big data analysis," *Science*, vol. 343, no. 6176, pp. 1203–1205, 2014.
- (29) R.A. Lewis, J.M. Rao, and D.H. Reiley, "Here, there, and everywhere: Correlated online behaviors can lead to overestimates of the effects of advertising," in *Proc. 20th Int. Conf. World Wide Web*. ACM, 2011, pp. 157–166.
- (30) D. Lovallo and O. Sibony, "The case of behavioral strategy," *McKinsey Quart.*, Mar. 2010.
- (31) A. McAfee, E. Brynjolfsson, T.H. Davenport, D.J. Patil, and D. Barton, "Big data: The Management Revolution," *Harvard Business Rev.*, vol. 90, no. 10, pp. 60–68, 2012.
- (32) E.M. Mirkes, T.J. Coats, J. Levesley, and A.N. Gorban, "Handling missing data in large healthcare dataset: A case study of unknown trauma outcomes," *Computers in Biology and Medicine*, vol. 75, pp. 203–216, 2016.
- (33) B.D. Mittelstadt and L. Floridi, "The ethics of big data: Current and foreseeable issues in biomedical contexts," *Science and Engineering Ethics*, vol. 22, no. 2, pp. 303–341, 2016.
- (34) J.M. Morse, "Cherry picking": Writing from thin data," *Qualitative Health Res.*, vol. 20, no. 1, p. 3, 2010.
- (35) Engineering National Academies of Sciences, Medicine *et al.*, *Refining the Concept of Scientific Inference When Working with Big Data: Proceedings of a Workshop*. National Acad. Press, 2017.
- (36) C. O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York, NY: Crown, 2016.
- (37) J. Pearl and D. Mackenzie, *The Book of Why: The New Science of Cause and Effect*, 1st ed. New York, NY: Basic, 2018.
- (38) M. Price and P. Ball, "Selection bias and the statistical patterns of mortality in conflict," *Statistical J. IAOS*, vol. 31, no. 2, pp. 263–272, 2015.
- (39) J. Qadir, A. Ali, R. ur Rasool, A. Zwitter, A. Sathiseelan, and J. Crowcroft, "Crisis analytics: Big data-driven crisis response," *J. of Int. Humanitarian Action*, vol. 1, no. 1, p. 12, 2016.
- (40) M.T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you?: Explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*. ACM, 2016, pp. 1135–1144.
- (41) J.C. Scott, *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven, CT: Yale Univ. Press, 1998.
- (42) T. Shelton, A. Poorthuis, and M. Zook, "Social media and the city: Rethinking urban socio-spatial inequality using user-generated geographic information," *Landscape and Urban Planning*, vol. 142, pp. 198–211, 2015.
- (43) N. Silver, *The Signal and the Noise: Why So Many Predictions Fail — But Some Don't*. Penguin, 2012.
- (44) U. Sivarajah, M.M. Kamal, Z. Irani, and V. Weerakkody, "Critical analysis of big data challenges and analytical methods," *J. Business Res.*, vol. 70, pp. 263–286, 2017.
- (45) G.D. Smith and S. Ebrahim, "Data dredging, bias, or confounding: They can all get you into the BMJ and the Friday papers," *BMJ: British Medical J.*, vol. 325, no. 7378, p. 1437, 2002.
- (46) L. Sweeney, "Simple demographics often identify people uniquely," *Health (San Francisco)*, vol. 671, pp. 1–34, 2000.
- (47) H.I. Weisberg, *Willful Ignorance*. Wiley, 2014.
- (48) E.W. Weisstein, Bonferroni correction, 2004; accessed Aug. 13, 2018.
- (49) H. White, "A reality check for data snooping," *Econometrica*, vol. 68, no. 5, pp. 1097–1126, 2000.
- (50) D.H. Wolpert, "The lack of a priori distinctions between learning algorithms," *Neural Computation*, vol. 8, no. 7, pp. 1341–1390, 1996.
- (51) R.V. Zicari and Andrej Zwitter, "Data for humanity: An open letter," *Scientific American*, Feb. 25, 2016; accessed Aug. 11, 2018.
- (52) A. Zwitter, *Humanitarian Intelligence: A Practitioner's Guide to Crisis Analysis and Project Design*. Rowman & Littlefield, 2016.
- (53) A. Zwitter and M. Boisse-Despiaux, "Blockchain for humanitarian action and development aid," *J. Int. Humanitarian Action*, vol. 3, no. 1, pp. 1–7, 2018.