

MIT Open Access Articles

*Fair, Transparent, and Accountable Algorithmic Decision-making
Processes*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Lepri, Bruno et al. "Fair, Transparent, and Accountable Algorithmic Decision-Making Processes." *Philosophy & Technology* 31, 4 (December 2018): 611–627 © 2017 Springer Science+Business Media

Published Version: <http://dx.doi.org/10.1007/s13347-017-0279-x>

Publisher: Springer Netherlands

Permanent Link: <https://hdl.handle.net/1721.1/122933>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: <http://creativecommons.org/licenses/by-nc-sa/4.0/>



Fair, transparent and accountable algorithmic decision-making processes

The premise, the proposed solutions, and the open challenges

Bruno Lepri · Nuria Oliver · Emmanuel Letouzé · Alex Pentland · Patrick Vinck

Received: date / Accepted: date

Abstract The combination of increased availability of large amounts of fine-grained human behavioral data and advances in machine learning is presiding over a growing reliance on algorithms to address complex societal problems. Algorithmic decision-making processes might lead to more objective and thus potentially fairer decisions than those made by humans who may be influenced by greed, prejudice, fatigue, or hunger. However, algorithmic decision-making has been criticized for its potential to enhance discrimination, information and power asymmetry, and opacity. In this paper we provide an overview of available technical solutions to enhance fairness, accountability and transparency in algorithmic decision-making. We also highlight the criticality and urgency to engage multi-disciplinary teams of researchers, practitioners, policy makers and citizens to co-develop, deploy and evaluate in the real-world algorithmic decision-making processes designed to maximize fairness and transparency. In doing so, we describe the Open Algorithms (OPAL) project as a step towards

Bruno Lepri
Fondazione Bruno Kessler, via Sommarive 18, Trento (Italy)
Tel.: +39-0461-314570
E-mail: lepri@fbk.eu

Nuria Oliver
Vodafone Research and Data-pop Alliance
E-mail: nuria.oliver@vodafone.com

Emmanuel Letouzé
Data-Pop Alliance and MIT Media Lab
E-mail: eletouze@datapopalliance.org

Alex Pentland
MIT and Data-pop Alliance
E-mail: pentland@mit.edu

Patrick Vinck
Harvard Humanitarian Initiative and Data-Pop Alliance
E-mail: pvinck@hsph.harvard.edu

realizing the vision of a world where data and algorithms are used as lenses and levers in support of democracy and development.

Keywords algorithmic decision-making · algorithmic transparency · fairness · accountability · social good

1 Introduction

Today’s vast and unprecedented availability of large-scale human behavioral data is profoundly changing the world we live in. Massive streams of data are available to train algorithms which, combined with increased analytical and technical capabilities, are enabling researchers, companies, governments and other public sector actors to resort to data-driven machine learning-based algorithms to tackle complex problems [26,67]. Many decisions with significant individual and societal implications previously made by humans alone –often by experts– are now made or assisted by algorithms, including hiring [12], lending [34], policing [66], criminal sentencing [5], and stock trading [33]. Data-driven algorithmic decision making may enhance overall government efficiency and public service delivery, by optimizing bureaucratic processes, providing real-time feedback and predicting outcomes [62]. In a recent book with the evocative and provocative title “Technocracy in America”, international relations expert Parag Khanna argued that a data-driven direct technocracy is a superior alternative to today’s (alleged) representative democracy, because it may dynamically capture the specific needs of the people while avoiding the distortions of elected representatives and corrupt middlemen [35]. Human decision making has often shown significant limitations and extreme bias in public policy, resulting in inefficient and/or unjust processes and outcomes [2, 23, 57, 65]. The turn towards data-driven algorithms can be seen as a reflection of a demand for greater objectivity, evidence-based decision-making, and a better understanding of our individual and collective behaviors and needs.

At the same time, scholars and activists have pointed to a range of social, ethical and legal issues associated with algorithmic decision-making, including bias and discrimination [4, 63], and lack of transparency and accountability [15, 47, 70, 69]. For example, Barocas and Selbst [4] showed that the use of algorithmic decision making processes could result in disproportionate adverse outcomes for disadvantaged groups, in ways suggestive of *discrimination*. Algorithmic decisions can reproduce and magnify patterns of discrimination, due to decision makers’ prejudices [46], or reflect the biases present in the society [46]. A recent study by ProPublica of the COMPAS Recidivism Algorithm (an algorithm used to inform criminal sentencing decisions by predicting recidivism) found that the algorithm was significantly more likely to label black defendants than white defendants, despite similar overall rates of prediction accuracy between the two groups [3]. Along this line, 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 asso-

ciated with big data-driven algorithmic decision-making, particularly in the areas of criminal justice and education [45].

In 2014, the White House released a report titled “Big Data: Seizing opportunities, preserving values” [50] highlighting the discriminatory potential of Big Data, including how it could undermine longstanding civil rights protections governing the use of personal information for credit, education, health, safety, employment, etc. For example, data-driven algorithmic 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 education will worsen their situation, and play against them in future applications. For this reason, a subsequent White House report called for “equal opportunity by design” as a guiding principle in those domains [43]. Furthermore, the White House Office of Science and Technology Policy, in partnership with Microsoft Research and others, has co-hosted several public symposiums on the impacts and challenges of algorithms and Artificial Intelligence, specifically relating to social inequality, labor, healthcare and ethics.¹

At the heart of the matter is the fact that technology outpaces policy in most cases; here, governance mechanisms of algorithms have not kept pace with technological development. Several researchers have recently argued that current control frameworks are not adequate for situations in which a potentially unfair or incorrect decision is made by a computer [4].

Fortunately, there is increasing awareness of the detrimental effects of discriminatory biases and opacity of some data-driven algorithmic decision-making systems, and of the need to reduce or eliminate them. A number of research and advocacy initiatives are worth noting, including the Data Transparency Lab², a “community of technologists, researchers, policymakers and industry representatives working to advance online personal data transparency through research and design”, and the DARPA Explainable Artificial Intelligence (XAI) project³. A tutorial on the subject was held at the 2016 ACM Knowledge and Data Discovery conference [27]. 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 these past few years on the ethical and legal challenges related to algorithmic governance and decision-making.⁴ Lepri *et al.* [38] recently discussed the need for *social good decision-making algorithms* (*i.e.* algorithms strongly influencing decision-making and resource optimization of public goods, such as public health, safety, access to finance and fair employment) to provide transparency and accountability, to only use personal information –created, owned and controlled by individuals–

¹ <https://www.whitehouse.gov/blog/2016/05/03/preparing-future-artificial-intelligence>

² <http://www.datatransparencylab.org/>

³ <http://www.darpa.mil/program/explainable-artificial-intelligence>

⁴ <http://www.law.nyu.edu/centers/ili/algorithmsconference>

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* by means of living lab approaches involving citizens.

In this paper, we focus on two of the main risks (namely, discrimination and lack of transparency) posed by data-driven predictive models leading to decisions that impact the daily lives of millions of people. There are additional challenges that we do not discuss in this paper. For example, issues relating to data ownership, privacy, informed consent and limited understanding (literacy) about algorithms' abilities and resulting risks among the general public are not discussed here. Instead, focusing on discrimination and lack of transparency, we provide the readers with a review of recent attempts at making algorithmic decision-making more fair and accountable, highlighting the merits and the limitations of these approaches. Finally, we turn to the description of a recent project, called Open Algorithms (OPAL), whose goal is to enable the design, implementation and monitoring of development policies and programs, accountability of government actions, and citizen engagement while leveraging the availability of large scale human behavioral data in a privacy-preserving and predictable manner.

2 Discriminatory effects of algorithmic decision-making

From a legal perspective, Tobler [64] described discrimination as “the application of different rules or practices to comparable situations, or of the same rule or practice to different situations”. Later, Barocas and Selbst [4] argued 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 [17]. Some have argued it could exacerbate prevailing inequalities by suggesting that historically disadvantaged groups actually deserve less favorable treatment [45].

Algorithmic discrimination may arise from different sources. 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 [14]. Second, discrimination can occur from the decision to use an algorithm itself. Categorization can be considered as a form of *direct discrimination*, whereby algorithms are used for disparate treatment [19]. Third, algorithms can lead to discrimination as a result of the misuse of certain models in different contexts [10]. 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 [10].

The use of algorithmic data-driven decision processes may also result in individuals 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 [51].

These are both old and new risks. There is ample evidence of detrimental impacts of current non-algorithmic approaches to access to finance, employment, and housing. Backgrounds checks for example are widely used in those procedures, with people's knowledge and consent. But hundreds of thousands of people have been treated unfairly as a result of mistakes (for instance, misidentification) in the procedures used by external companies to perform background checks⁵. On the one hand, the occurrence of such trivial procedural mistakes may be bound to decrease once performed through data-driven methodologies. But 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. But it also means that should such methodologies not be transparent in their inner workings, the effects are likely to stay though with different roots.

This scenario thus highlights particularly well the need for machine learning models featuring transparency (understood as openness and communication of both the data being analyzed and the mechanisms underlying the models), accountability (understood as the assumption of accepting the responsibility for actions and decisions) and fairness (understood as the lack of discrimination or bias in the decisions).

3 Techniques to prevent algorithmic discrimination and maximize fairness

A simple way to try to maximize *fairness* –understood as the lack of bias and discrimination– in machine learning is precluding the use of sensitive attributes [9, 31, 60, 4]. For example, if we want a race-blind or a gender-blind decision-making process we may exclude these attributes (*i.e.* race, gender, etc.) from the process. However, this solution has several technical problems. First, the excluded attributes can often be implicit in non-excluded ones [48, 56, 70]. For example, when race is excluded as a criterion for granting or not a loan, some implicit information can be present in the individual's zip code, given that zip code may be a good proxy for race [60, 41]. An additional problem with the *blindness approach* was identified by Dwork *et al.* [20]: a learning decision rule can select the opposite of what is intended. Consider the example provided by Dwork *et al.* [20]: in a certain culture S the most talented students tend to

⁵ <http://www.chicagotribune.com/business/ct-background-check-penalties-1030-biz-20151029-story.html>

study engineering and science whereas the less talented study finance. However, in another culture C the trend is reversed such that the most talented students are encouraged to study finance and the less talented are guided towards engineering and science. An organization from culture C ignorant of these cultural differences, may select candidates for “economics”, potentially selecting the wrong candidates from culture S even while maintaining parity. This is an example of a suboptimal outcome in a “fairness through blindness” approach as the errors are due to ignoring cultural membership.

In the last few years, several researchers have proposed different technical definitions of *fairness* in machine learning, most of which formalize some notion of *group fairness* [9, 32, 71, 22]. One of the most used notions is *statistical parity*, which requires that an equal fraction of each group should receive each possible outcome [9, 32, 71, 22]. Recent papers have also considered approximate relaxations of statistical parity, motivated by the formulation of disparate impact in the U.S. legal code [22, 68]. Work in these directions has also developed learning algorithms that penalize violations of statistical parity [9, 32].

However, as pointed out by Dwork *et al.* [20], the *group fairness* often fails at both accurate learning and actual fairness. The case of lending can be used as an example to make the point: if two groups have different proportions of individuals who are able to pay back their loans, the algorithm’s accuracy will suffer when constrained to predict an equal proportions of paybacks for the two groups. Moreover, *group fairness* definitions do not guarantee that a creditworthy individual from one group has an equal probability of receiving a loan as a similarly creditworthy individual from the other group.

To overcome these limitations, Dwork *et al.* [20] argued that technical definitions of fairness should focus on *individual fairness*. More precisely, they proposed a framework based on a task-specific externally defined similarity metric between individuals. The goal of this metric is to achieve fairness through the principle that similar people should be treated in a similar way: thus, any two individuals who are similar with respect to a given task should be classified in a similar way [20]. The technical definition of similarity between individuals, proposed by Dwork *et al.* [20], resembles partly the notion of *strict equality of opportunity* proposed by the political scientist John Roemer [54, 55]. For Roemer, *strict equality of opportunity* is achieved when people, irrespectively of circumstances beyond their control, have the same ability to achieve advantage through their free choices.

Following Dwork *et al.* [20], Joseph *et al.* [30] have recently proposed a specific definition of individual fairness that can be considered as a mathematical formalization of the Rawlsian principle of “fair equality of opportunity” [52]. This principle affirms that those individuals, “who are at the same level of talent and have the same willingness of using it, should have the same perspectives of success regardless their initial place in the social system” (*e.g.* income, race, etc.) [52]. Thus, this principle is stronger than a “formal equality of opportunity”: Rawls, indeed, argued that an individual should not only have the right to opportunities, but also should have an effective equal chance

as another individual of similar natural abilities. In their proposed approach, Joseph *et al.* [30] include a notion of fairness in a sequential decision-making framework called contextual bandits in the machine learning literature. Their notion of fairness requires that at every step the learning algorithm never favors applicants whose attributes are lower than the ones of another applicant. Hence, their aim is to design a machine learning algorithm that would (provably) converge to an optimal decision while being (provably) fair at every step. They show that learning algorithms can be proven to be fair in such a way that the cost (from the perspective of rate of convergence to an optimal decision) of adding fairness to the algorithm is small.

In a recent work, Hardt *et al.* [29] proposed a fairness measure, based on a similar notion of *equality of opportunity*, that tries to achieve two important objectives. First, to overcome the main conceptual shortcomings of statistical parity as a fairness notion. Second, to build higher accuracy classifiers, in line with the central goal of supervised machine learning. To this end, they proposed a criterion for discrimination against a specified sensitive attribute in supervised learning, where the goal is to predict some target based on available features. Assuming the availability of data about the predictor, the target, and the membership in the protected group, they showed how to optimally adjust any learned predictor so as to remove discrimination according to their definition. The proposed framework also changes incentives by shifting the cost of poor classification from disadvantaged groups to the decision maker, who can respond by improving the classification accuracy. They illustrate their approach and compare different fairness measures in the case of FICO scores with the protected attribute of race, *i.e.* they build machine learning models that aim to predict a credit risk from a number of attributes with race being a protected attribute. In their conclusions, they highlight that with their framework it is possible to measure unfairness rather than to prove fairness and emphasize the importance of having access to reliable target variables, which is not always the case in practical scenarios.

Another interesting result is the one discussed by Kleinberg *et al.* [36]. In their paper, they formalized three fairness conditions that constitute the heart of the debates about discrimination in machine learning. Moreover, they proved that, except in highly constrained special cases, there is no method that can satisfy these three conditions simultaneously. Specifically, a first condition—known as *calibration within groups* in the literature—is that the probability estimates provided by the decision-making algorithms should be well-calibrated: for example, if the algorithm identifies a set of people as having probability z of constituting positive instances, then approximately a z fraction of this set should be positive instances [24]. In addition, this condition should be valid when applied separately in each group: if we think of potential differences between an outcome z for Afro-Americans and Asians, this means that a z fraction of men and z fraction of women assigned a probability z should possess the property in question. A second condition focuses on the people who constitute positive instances: the average score received by those people should be the same in each group. This represents a *balance for the*

positive class: indeed a violation of this condition would mean that people who are positive instances in one group receive consistently lower probability estimates than people constituting positive instances in another group. Let resort to the case study investigated by ProPublica where one of the concerns raised was that white defendants who went on to commit future crimes were assigned risk scores corresponding to lower probability estimates in aggregate. This is an example of a violation of the *balance for the positive class* condition. A similar condition holds with respect to negative instances, which is called *balance for the negative class*. In short, these balance conditions can be considered as generalizations of the notions that both groups should have equal false negative and false positive rates. The authors outline a few lines of future research, including the fact that there might be use cases where the cost of false positives differs greatly from the cost of false negatives and thus it should be taken into account. In line with this approach, Chouldechova [13] and Corbett-Davies *et al.* [16] consider conditions close to the balance for negative and positive classes together with a form of calibration adapted to binary predictions. The calibration requires that for all people given a positive label, the same fraction of people in each group should truly be part of the positive class. Interestingly, they show that no classification rule can satisfy the required constraints. Finally, a paper by Friedler *et al.* [25] defines two axiomatic properties of feature generation and shows that no mechanism can be fair under these two properties.

The results obtained by Kleinberg *et al.* [36], Chouldechova [13], and Corbett-Davies *et al.* [16] highlight that it is not enough to simply demand *algorithmic fairness*. We may need to investigate deeply and critically each problem and determine which notion of fairness is considered to be the most relevant and meaningful for that particular problem. Critically, what constitutes fairness changes according to different worldviews: for example, the philosopher Robert Nozick in his book “Anarchy, State, and Utopia” [44] proposed a libertarian alternative view where he is concerned that eliminating discrimination biases, present in society, may create new harms to new groups of people. Instead, Dworkin’s egalitarian view of fairness [21] is based on the principle of *equality of resources*.

For this reason, we feel the urgency to establish a call for action putting together researchers from different fields –including law, ethics, political philosophy and machine learning– to devise, evaluate and validate in the real-world alternative fairness metrics for different tasks. In addition to this empirical research, we believe it will be necessary to propose a modeling framework – supported by empirical evidence– that would assist practitioners and policy makers in making decisions aided by algorithms that are maximally fair.

4 Information asymmetry and lack of transparency

The mandate for accountable algorithms in government and corporations’ decision-making tools is fundamental in both validating their utility toward

the public interest as well as redressing potential harms generated by these algorithms.

Transparency, which refers to the understandability of a specific model, can be a mechanism that facilitates accountability. More specifically, transparency can be considered at the level of the entire model, at the level of individual components (*e.g.* parameters), and at the level of a particular training algorithm. In the strictest sense, a model is transparent if a person can contemplate the entire model at once. Thus, models should be characterized by low computational complexity. A second and less strict notion of transparency might be that each part of the model (*e.g.* each input, parameter, and computation) admits an intuitive explanation [40]. A final notion of transparency might apply at the level of the algorithm, even without the ability to simulate an entire model or to intuit the meaning of its components.

However, the ability to access and analyze behavioral data about customers and citizens on an unprecedented scale gives corporations and governments powerful means to reach and influence segments of the population through targeted marketing campaigns and social control strategies. In particular, we are witnessing an *information asymmetry* situation where a powerful few have access and use resources and tools that the majority do not have access to, thus leading to an –or exacerbating the existing– asymmetry of power between the state and big companies on one side and the people on the other side [1], conceptualized as a “new digital divide” [7]. In addition, the nature and use of various data-driven algorithms for social good, as well as the lack of computational or data literacy among citizens [6], makes algorithmic transparency difficult to generalize and accountability difficult to assess [47].

Burrell [8] 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 legislation⁶. But powerful commercial and governmental interests will make it difficult to eliminate intentional 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 “algorithmic literacy” and by enabling independent experts to advise those affected by algorithmic 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

⁶ 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)

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.

5 Techniques to improve transparency and accountability

As previously described, algorithmic decision-making might lack transparency. A simple solution to this limitation would consist of asking for transparency and openness of the algorithm’s source code as well as inputs and outputs that are used to make relevant algorithmic decisions. However, transparency alone is not sufficient to provide accountability in all cases. First of all, it is often necessary to keep secret certain elements of an algorithmic decision policy, the way how the policy is implemented, the key inputs, or the outcome. This is a way to help prevent strategic *gaming* of the system. Furthermore, when the decision being regulated is a commercial one –such as a bank decision to give a loan–, a legitimate business interest in protecting proprietary information or algorithms may be incompatible with full transparency. Again, an algorithmic decision-making system may use as input or may create as output sensitive data that should be not shared to protect business interests, privacy, etc. In some domains, such as finance and healthcare, disclosure may be limited by regulations.

A strategy for verifying and making transparent the behavior of an algorithmic decision-making process is *auditing*. An auditing strategy deals with the decision process as a black box, whose inputs and outputs are visible, while inner workings are not [59]. However, several researchers have shown that a black-box evaluation of decision processes and systems is the least powerful of a set of available methods for understanding their behaviors [18]. Datta *et al.* [18] study the *opacity* (*i.e.* lack of transparency) of web-based ads by means of their AdFisher tool. They ran several experiments to investigate the transparency provided by Google’s Ad Settings. In particular, they analyze if visiting webpages related to a certain interest would lead to a change in the ads shown that is not captured in the settings. They found instances of opacity as they encountered cases where there were significant differences in the ads shown to different profiles while their tool failed to show any type of profiling. They attribute the instances of opacity to remarketing or to other causes related to the complexity of Google’s massive, automated advertising system. The authors advocate –as we do– for additional research to create machine learning algorithms that automatically provide transparency to the users.

Furthermore, some algorithmic decision-making processes aim to determine variables that are not measurable in a direct way, for example the risk of credit default. For this reason, these values are computed by means of proxy variables such as the consumer’s credit history, the consumer’s income, some consumer’s personal characteristics or even their mobile phone usage patterns [58]. Consumers able to understand these processes would be tempted to con-

trol the proxy variables [45]. Thus, secrecy discourages consumers' strategic behaviors and prevents violations of legal restrictions on disclosure of data. In a recent paper, Hardt *et al.* [28] have proposed adversarial methods for designing algorithmic decision-making processes that remain robust in the face of gaming. Moreover, for algorithmic decision processes involving some element of randomness, the full transparency of the source code, inputs, operating environment, and results does not exclude the possibility that the process may produce unpredictable results. Finally, systems that change over time cannot be fully understood through transparency alone. For example, online machine learning algorithms can update their model for predictions after each decision, which increases the complexity of a strategy to ensure transparency in the decision-making process.

Ultimately, we need accountability in decision-making algorithms such that there is clarity regarding who holds the responsibility of the decisions made by them or with algorithmic support. Transparency is generally thought as a key enabler of accountability. However, transparency and auditing do not necessarily suffice for accountability. In fact, in a recent paper Kroll *et al.* [37] have introduced computational methods able to provide accountability even when some information is kept hidden. The authors used advanced techniques to allow the governance of secret algorithmic decision-making processes: specifically, software verification (*i.e.* a set of techniques for proving mathematically that a piece of software has certain properties), cryptographic commitments (*i.e.* equivalents of sealed documents held by a third party or in a safe place), zero-knowledge proofs (*i.e.* cryptographic tools that allow a decision maker, as part of a cryptographic commitment, to prove that the decision policy that was actually used has a certain property, but without revealing either how the property is known or what the decision policy is), and fair random choices (*i.e.* a technique allowing software that makes random choices to be fully reproducible). These methods can guarantee that both the input data and the software that analyzes such data satisfy the requirements for procedural regularity, even when they are kept secret.

Another approach to provide transparency in algorithmic decision-making entails providing explanations regarding the processes that lead to the decisions such that they are interpretable by humans. In a recent work, Ribeiro *et al.* [53] proposed to (i) provide explanations for individual predictions as a solution to the so-called "trusting a prediction" problem, and (ii) select multiple such predictions and explanations as a solution to the so-called "trusting the model" problem. Specifically, they proposed LIME, a novel technique that explains the predictions of any classifier by learning an interpretable model locally around the prediction. They also proposed a method to explain models by showing representative individual predictions and their explanations in a non-redundant way. In the paper, they showed the value of these explanations by means of experiments, both simulated and with human subjects, on several scenarios, such as (i) deciding if one should trust a prediction, (ii) identifying a classifier that should not be trusted, and (iii) choosing between different classification models, etc.

A recent challenging contribution by Lipton [39] examined the motivations underlying the raising interest in *interpretability*, finding them to be diverse and sometimes discordant. In general, the desire for an interpretation suggests that predictions alone are not sufficient, thus implying a discrepancy between the real-world objectives of machine learning researchers and practitioners and the simple objectives optimized by most machine learning models.

However, the concept of *interpretability* is often vaguely defined. The learning model’s properties that enable or that compromise interpretability broadly fall into two categories. The first one relates to *transparency*, that is *how does the model work*. The second one consists of *post-hoc interpretations*, that is *what else can the model tell*.

As distinct notion of interpretability, post-hoc interpretations consist of explanations that need not elucidate the exact process by which models work. These interpretations include natural language explanations [42], visualizations of learned representations or models (*e.g.* saliency maps in deep neural nets [61]), and explanations by example (*e.g.* a tumor is classified as malignant because to the model it looks like these other tumors) [11]. One advantage of this concept of interpretability is that we can interpret opaque models after-the-fact, without sacrificing predictive performance.

Algorithmic decision-making processes have the potential to lead to fairer and more objective decisions, grounded in data that are representative of the community where the decisions apply. However, as explained in the previous sections, algorithmic decision-making might lead to discrimination, information asymmetry and lack of transparency. Hence, we believe that it is not only important but also urgent to engage multi-disciplinary teams of researchers, practitioners and policy makers to propose, implement and evaluate in the real-world algorithmic decision-making processes that are designed to maximize their fairness and transparency.

In the next section, we describe one of such proposals, the OPAL project.

6 The OPAL project

The Open Algorithms (OPAL) project⁷, is a multi-partner socio-technological platform led by Data-Pop Alliance, Imperial College London, the MIT Media Lab, Orange S.A., and the World Economic Forum, that aims to leverage private sector data for public good purposes by “sending the code to the data” in a privacy preserving, predictable, participatory, scalable and sustainable manner. The project came out of the recognition that accessing data held by private companies –including Call Detail Records (CDRs) collected by telecom operators for billing purposes, and banking data– for research and policy purposes has been a conundrum. To date for example, CDRs have been accessed and analyzed either internally, or externally through ad-hoc data challenges or through bilateral arrangements with a limited number of groups under

⁷ <http://opalproject.org/>

Non-Disclosure Agreements. These types of engagements have offered ample evidence of the promise and demand, but they do not scale nor address some of the most critical challenges discussed above.

Building on the lessons of the past, OPAL is a key milestone towards a vision where data is at the heart of societal development around the globe, by providing a far better picture of human conditions to official statisticians, policy makers, planners, businesses leaders, and citizens, while enabling greater inclusion and inputs of all members of societies on the kinds and uses of analyses performed on data about themselves. As such, OPAL will reflect and foster the double objective to turn Big Data on its head and “save it from itself” [49].

OPAL’s first core feature is technological: the platform allows sending queries (i.e. running algorithms) on the partner companies’ servers, behind their firewalls, and not the other way around, so that raw data are never exposed to theft and misuse. The second one is socio-political: it involves co-designing the algorithms so that they are not only open but also serve local needs and respect local standards. For example, a key idea of OPAL is that algorithms should be verified by experts, policy makers, citizens to be as free as possible from biases and unintended side effects such as discrimination. To this end, OPAL makes use of vetting the algorithms that are permitted to run on a given data-set within a specific data repository. Once an algorithm has been vetted, it becomes a template that is digitally signed by the issuers (e.g. expert themselves, public institutions, representatives of the communities that will be affected by algorithmic decisions, etc.). This template algorithm can be shared among a group of entities (e.g. within a consortium) or even be published on a public site. Note that this vetting does not guarantee the quality of the output, which is a function of the quality of the input data.

The OPAL model of moving the algorithm to the data and of using vetting allows a data repository to choose whether or not it is willing to accept a submitted OPAL algorithm (query). If the data repository accepts a given vetted algorithm, it also has the option to impose additional filtering on the resulting data prior to being returned as answer to the querier (e.g. defining the degree of personal information within a given answer). Moreover, each repository can introduce machine learning algorithms as additional mechanisms to protect privacy. Such algorithms allow a repository to detect if multiple accesses from the same entity may result in compromising Personal Identifiable Information (PII).

Finally, blockchain technology can be used to capture and log both vetted-queries and safe-answers, thus providing a mechanism to support post-event audit and accountability. One easy way would be for the querier to compute a cryptographic hash of the query sent, and for the data repository to compute the hash of the response. In addition, the technology may be used by data owners when they want to monetize their data, including the ability to link money and data flows, or micro-payments with small transaction costs.

OPAL is currently being deployed through pilots in Senegal and Colombia, where it has been endorsed by and benefits from the support of their National

Statistical Offices and major local telecom operators. Local engagement and empowerment will be central to the development of OPAL: needs, feedback and priorities have been collected and identified through local workshops and discussions, and their results will feed into the design of future algorithms. These algorithms will be fully open, therefore subject to public scrutiny and redress. A local advisory board is being set up to provide guidance and oversight to the project. In addition, trainings and dialogues will be organized around the project to foster its use and diffusion as well as local capacities and awareness more broadly. OPAL aims to be deployed in 2 more countries by the end of 2018.

Initiatives such as OPAL have the potential to enable more human-centric accountable and transparent data-driven decision-making and governance, by involving a wide range of stakeholders in and through their design and implementation. Note, however, that OPAL does not fully address the issues of algorithmic fairness and transparency. It does not address internal uses of data and potentially discriminatory behavior by corporations. It does not advance an open data agenda, which may be unrealistic when working with corporate-owned data carrying high competitive value and posing severe privacy risks. It also does not in itself enable control for bias in the data itself.

Despite these limitations, we believe that OPAL will provide an avenue for an array of users, from official statisticians to community organizers, to openly query data and have those queries and results examined through a fairness and anti-discrimination lens. It is an important concrete step towards a world where data and algorithms can be leveraged through participatory processes for societal development and democracy around the globe.

7 Conclusions

We live in an unprecedented historic moment where the availability of vast amounts of human behavioral data, combined with advances in machine learning are enabling us to tackle complex problems through algorithmic decision-making. The opportunities to have positive social impact through fairer and more transparent decisions are paramount. However, algorithmic decision-making processes might lead to discrimination, information asymmetry and lack of transparency.

In this paper we have provided an overview of both existing limitations and proposed solutions regarding fairness, accountability and transparency in algorithmic decision-making. We have highlighted open challenges that would still need to be addressed and have described the OPAL project as an exemplary effort that aims to maximize algorithmic fairness and transparency to support decision-making for social good. We would like to emphasize the importance and the urgency to engage multi-disciplinary teams of researchers, practitioners and policy makers to propose, implement and evaluate in the real-world algorithmic decision-making processes that are designed to maximize their fairness and transparency.

The opportunity to significantly improve the processes leading to decisions that affect millions of lives is huge. As researchers and citizens we believe that we should not miss on this opportunity. Hence, we would like to encourage the larger community –*e.g.* researchers, practitioners, policy makers– in a variety of fields –*e.g.* computer science, sociology, economics, ethics, law– to join forces so we can address today’s limitations in data-driven decision-making and contribute to fairer and more transparent decisions with clear accountability and significant positive impact.

References

1. Akerlof, G.: The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics* **84**(3), 488–500 (1970)
2. Akerlof, G., Shiller, R.: *Animal spirits: How human psychology drives the economy, and why it matters for global capitalism*. Princeton University Press (2009)
3. Angwin, J., Larson, J., Mattu, S., Kirchner, L.: Machine bias. *ProPublica* (2016). URL <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
4. Barocas, S., Selbst, A.: Big data’s disparate impact. *California Law Review* **104**, 671–732 (2016)
5. Barry-Jester, A.M., Casselman, B., Goldstein, D.: The new science of sentencing. *The Marshall Project* (2015). URL <https://www.themarshallproject.org/2015/08/04/the-new-science-of-sentencing>.
6. Bhargava, R., Deahl, E., Letouzé, E., Noonan, A., Sangokoya, D., Shoup, N.: Beyond data literacy: Reinventing community engagement and empowerment in the age of data. *Data-Pop Alliance White Paper Series* (2015). URL <http://datapopalliance.org/wp-content/uploads/2015/11/Beyond-Data-Literacy-2015.pdf>
7. boyd, d., Crawford, K.: Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, Communication, & Society* **15**(5), 662–679 (2012)
8. Burrell, J.: How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society* **3**(1) (2016)
9. Calders, T., Verwer, S.: Three naive bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery* **21**(2), 277–292 (2010)
10. Calders, T., Zliobaite, I.: Why unbiased computational processes can lead to discriminative decision procedures. In: B. Custers, T. Calders, B. Schermer, T. Zarsky (eds.) *Discrimination and Privacy in the Information Society*, pp. 43–57 (2013)
11. Caruana, R., Kangaroo, H., David, J., Dionisio, N., Sinha, U., Johnson, D.: Case-based explanation of non-case-based learning methods. In: *Proceedings of the 1999 American Medical Informatics Association (AMIA) Symposium*, pp. 212–215 (1999)
12. Chalfin, A., Danieli, O., Hillis, A., Jelveh, Z., Luca, M., Ludwig, J., Mullainathan, S.: Productivity and selection of human capital with machine learning. *American Economic Review* **106**(5), 124–127 (2016)
13. Chouldechova, S.: Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *arXiv preprint arXiv:1610.07524* (2016)
14. Christin, A., Rosenblatt, A., boyd, d.: Courts and predictive algorithms. *Data & Civil Rights Primer* (2015)
15. Citron, D., Pasquale, F.: The scored society. *Washington Law Review* **89**(1), 1–33 (2014)
16. Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., Huq, A.: Fair algorithms and the equal treatment principle. *Working Paper* (2017)
17. Crawford, K., Schultz, J.: Big data and due process: Toward a framework to redress predictive privacy harms. *Boston College Law Review* **55**(1), 93–128 (2014)
18. Datta, A., Tschantz, M.C., Datta, A.: Automated experiments on ad privacy settings. In: *Proceedings on Privacy Enhancing Technologies*, pp. 92–112 (2015)

19. Diakopoulos, N.: Algorithmic accountability: Journalistic investigation of computational power structures. *Digital Journalism* (2015)
20. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.: Fairness through awareness. In: *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pp. 214–226. ACM (2012)
21. Dworkin, R.: *Sovereign virtue: The theory and the practice of equality*. Harvard University Press (2000)
22. Feldman, M., Friedler, S., Moeller, J., Scheidegger, C., Venkatasubramanian, S.: Certifying and removing disparate impact. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 259–268 (2015)
23. Fiske, S.: Stereotyping, prejudice, and discrimination. In: D. Gilbert, S. Fiske, G. Lindzey (eds.) *Handbook of Social Psychology*, pp. 357–411. Boston: McGraw-Hill (1998)
24. Foster, D., Vohra, R.V.: Asymptotic calibration. *Biometrika* **85**(2), 379–390 (1998)
25. Friedler, S.A., Scheidegger, C., Venkatasubramanian, S.: On the (im)possibility of fairness. arXiv preprint [arXiv:1609.07236](https://arxiv.org/abs/1609.07236) (2016)
26. Gillespie, T.: The relevance of algorithms. In: T. Gillespie, P. Boczkowski, K. Foot (eds.) *Media technologies: Essays on communication, materiality, and society*, pp. 167–193. MIT Press (2014)
27. Hajian, S., Bonchi, F., Castillo, C.: 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*, pp. 2125–2126. ACM (2016)
28. Hardt, M., Megiddo, N., Papadimitriou, C., Wootters, M.: Strategic classification. In: *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pp. 111–122. ACM (2016)
29. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. In: *Proceedings of the International on Advances in Neural Information Processing Systems (NIPS)*, pp. 3315–3323 (2016)
30. Joseph, M., Kearns, M., Morgenstern, J., Neel, S., Roth, A.: Rawlsian fairness for machine learning. arXiv preprint [arXiv:1610.09559](https://arxiv.org/abs/1610.09559) (2016)
31. Kamiran, F., Calders, T., Pechenizkiy, M.: Discrimination aware decision tree learning. In: *Proceedings of 2010 IEEE International Conference on Data Mining*, pp. 869–874. IEEE (2010)
32. Kamishima, T., Akaho, S., Asoh, H., Sakuma, J.: Fairness-aware classifier with prejudice remover regularizer. In: *Proceedings of the European Conference on Machine Learning and Principles of Knowledge Discovery in Databases (ECMLPKDD), Part II*, pp. 35–50 (2011)
33. Kearns, M., Nevmyvaka, Y.: Machine learning for market microstructure and high frequency trading. In: M. O’Hara, M. Lopez de Prado, D. Easley (eds.) *High Frequency Trading*. Risk Books (2013)
34. Khandani, A.E., Kim, A.J., Lo, A.W.: Consumer credit risk models via machine-learning algorithms. *Journal of Banking and Finance* **34**, 2767–2787 (2010)
35. Khanna, P.: *Technocracy in America: Rise of the info-state*. CreateSpace Independent Publishing Platform (2017)
36. Kleinberg, J., Mullainathan, S., Raghavan, M.: Inherent trade-offs in the fair determination of risk scores. In: *Proceedings of the 8th Innovations in Theoretical Computer Science Conference*. ACM (2017)
37. Kroll, J.A., Huey, J., Barocas, S., Felten, E.W., Reidenberg, J.R., Robinson, D.G., Yu, H.: Accountable algorithms. *University of Pennsylvania Law Review* **165** (2017)
38. Lepri, B., Staiano, J., Sangokoya, D., Letouz, E., Oliver, N.: The tyranny of data? the bright and dark sides of data-driven decision-making for social good. arXiv preprint [arXiv:1612.00323](https://arxiv.org/abs/1612.00323) (2017)
39. Lipton, Z.C.: The mythos of model interpretability. In: *2016 ICML Workshop on Human Interpretability in Machine Learning* (2016)
40. Lou, Y., Caruana, R., Gehrke, J., Hooker, G.: Accurate intelligible models with pairwise interactions. In: *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 623–631. ACM (2012)
41. Macnish, K.: Unblinking eyes: The ethics of automating surveillance. *Ethics and Information Technology* **14**(2), 151–167 (2012)

42. McAuley, J., Leskovec, J.: Hidden factors and hidden topics: Understanding rating dimensions with review text. In: Proceedings of the 7th ACM conference on Recommender Systems (2013)
43. Munoz, C., Smith, M., Patil, D.: Big data: A report on algorithmic systems, opportunity, and civil rights. Tech. rep., Executive Office of the President (2016)
44. Nozick, R.: Anarchy, state, and utopia. Basic Books (1974)
45. O’Neil, C.: Weapons of math destruction: How big data increases inequality and threatens democracy. Crown (2016)
46. Pager, D., Shepherd, H.: The sociology of discrimination: Racial discrimination in employment, housing, credit and consumer market. *Annual Review of Sociology* **34**, 181–209 (2008)
47. Pasquale, F.: The Black Box Society: The secret algorithms that control money and information. Harvard University Press (2015)
48. Pedreschi, D., Ruggieri, S., Turini, F.: Discrimination-aware data mining. In: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 560–568 (2008)
49. Pentland, A.: Saving big data from itself. *Scientific American* **311**(2), 64–67 (2014)
50. Podesta, J., Pritzker, P., Moniz, E., Holdren, J., Zients, J.: Big data: Seizing opportunities, preserving values. Tech. rep., Executive Office of the President (2014)
51. Ramirez, E., Brill, J., Ohlhausen, M., McSweeney, T.: Big data: A tool for inclusion or exclusion? Tech. rep., Federal Trade Commission (2016)
52. Rawls, J.: A theory of justice. Harvard University Press (1971)
53. Ribeiro, M., Singh, S., Guestrin, C.: “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, pp. 1135–1144 (2016)
54. Roemer, J.E.: Theories of Distributive Justice. Harvard University Press (1996)
55. Roemer, J.E.: Equality of Opportunity. Harvard University Press (1998)
56. Romei, A., Ruggieri, S.: A multidisciplinary survey on discrimination analysis. *The Knowledge Engineering Review* **29**(5), 582–638 (2014)
57. Samuelson, W., Zeckhauser, R.: Status quo bias in decision making. *Journal of Risk and Uncertainty* (1), 7–59 (1988)
58. San Pedro, J., Proserpio, D., Oliver, N.: Mobiscore: Towards universal credit scoring from mobile phone data. In: Proceedings of the International Conference on User Modeling, Adaptation and Personalization (UMAP), pp. 195–207 (2015)
59. Sandvig, C., Hamilton, K., Karahalios, K., Langbort, C.: Auditing algorithms: Research methods for detecting discrimination on internet platforms. In: Data and Discrimination: Converting Critical Concerns into Productive Inquiry, a preconference at the 64th Annual Meeting of the International Communication Association (2014)
60. Schermer, B.W.: The limits of privacy in automated profiling and data mining. *Computer Law & Security Review* **27**(1), 45–52 (2011)
61. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: Visualising image classification models and saliency maps. arXiv preprint **arXiv:1312.6034** (2013)
62. Sunstein, C.: Regulation in an uncertain world. National Academy of Sciences (2012). URL <https://www.whitehouse.gov/sites/default/files/omb/inforeg/speeches/regulation-in-an-uncertain-world-06202012.pdf>
63. Sweeney, L.: Discrimination in online ad delivery. Available at SSRN: <http://ssrn.com/abstract=2208240> (2013)
64. Tobler, C.: Limits and potential of the concept of indirect discrimination. Tech. rep., European Network of Legal Experts in Anti-Discrimination (2008)
65. Tversky, A., Kahnemann, D.: Judgment under uncertainty: Heuristics and biases. *Science* **185**(4157), 1124–1131 (1974)
66. Wang, T., Rudin, C., Wagner, D., Sevieri, R.: Learning to detect patterns of crime. In: Machine Learning and Knowledge Discovery in Databases, pp. 515–530. Springer (2013)
67. Willson, M.: Algorithms (and the) everyday. *Information, Communication & Society* (2016)
68. Zafar, M.B., Martinez, I.V., Rodriguez, M.D., Gummadi, K.P.: Learning fair classifiers. arXiv preprint **arXiv:1507.05259** (2015)

-
69. Zarsky, T.: Automated prediction: Perception, law and policy. *Communications of the ACM* **4**, 167–186 (1989)
 70. Zarsky, T.: 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)
 71. Zemel, R., Wu, Y., Swersky, K., Pitassi, T., Dwork, C.: Learning fair representation. In: *Proceedings of the 2013 International Conference on Machine Learning (ICML)*, pp. 325–333 (2012)