Mozilla's comments to the UK Algorithms Inquiry

18 October 2017

Introduction

1. Mozilla is a global community working together to build a better Internet. As a mission-driven organization, we are dedicated to promoting openness, innovation, and opportunity online. As a core principle, we believe that the Internet, as the most significant social and technological development of our time, is a precious public resource that must be improved and protected.

2. Algorithms are the heart of computing, but have grown to drive many devices far beyond a typical computer. The concept encompasses many different kinds of programming, decision-making, and learning processes. These are powerful tools that can be used to pursue any set of goals: whether in the interest of society or a particular actor or set of actors. They can also insert bias into decisions, whether inadvertent or in pursuit of those goals. With appropriate consideration in creation, monitoring, and governance of these tools (including algorithms, machine learning, and data sets) we can shape these for social good. Through this submission, we hope to help policymakers understand the breadth and variety of automated decision-making and the importance of the data sets both for machine learning and testing the impact of algorithms.

Executive Summary

3. This is a new and evolving area in which even the terms we use to discuss the topic are not fixed. The committee should spend time working out how to frame good questions, as well as looking for answers.

4. "Algorithm" is too general a concept to reason about. As well as having different uses, algorithms can be of several different types, and operated by either government or private sector organizations; those factors should have a significant effect on how we view them. Furthermore, the context of the decision-making is, in many cases, at least as important than the algorithm itself.

5. Algorithms and data are inherently increasingly intertwined when there is machine learning; the data shapes the decision-making process. Just talking about one without the other will give an incomplete picture, and yet ideas of "data transparency" and "algorithm transparency" have very different issues and challenges.

Inquiry

Algorithms

6. Thank you for the opportunity to submit evidence to the committee on this topic. The relationship of algorithms to people and society, and the potentially negative effects that misapplied or miscalibrated algorithms can have on them, is a new and evolving area in which even the terms we use to discuss the topic are not fixed. We hope our and other

submissions will help the committee understand the context of this discussion as much as providing answers to questions.

7. We suggest that “algorithm” is too general a word to reason about without clearer definition. We see at least 3 types of algorithm in use today:

    a. **Operational**: Algorithms used in operations - using code to automate a manual process. While such systems are generally not what people conceptualize when talking about fairness of algorithms, we include this category merely to state that these too are “algorithms,” though of a sort that does not by itself require special regulation.

    b. **Scoring**: Scoring and threshold-based decision-making, usually about a person and often used to determine eligibility for credit, loans, insurance, or employment.

    c. **Presentation**: Prioritizing, presenting, and/or personalizing information, such as in social media, news feeds, and search results.

8. The impact of algorithms in the “Scoring” class has been recognised for a longer time, and algorithms in that class have been regulated before on a per-industry basis - e.g. in the financial sector in the EU, or by the Fair Credit Reporting Act in the USA. But the use of both Scoring and Presentation algorithms can have significant effect on a person’s life, options or even opinions, and there are fewer laws directed specifically to Presentation-type algorithms.

9. We suggest that algorithms in the Operations class should not, in our view, be subject to regulatory oversight. They are relatively straightforward, and generally are not distinct from pre-automation functionality. As such, we will set those aside for the purposes of this inquiry.

10. Scoring and Presentation algorithms, in contrast, are often highly sophisticated, and have a wide-range of context-dependent inputs that are not clear to the end user. Similarly, their execution and outputs may depend on a decision-making or scoring process that is complex or evolving. These factors make straightforward seeming concepts of transparency and fairness very difficult to achieve in practice, as important as the principles behind these goals are. Consequently, context is particularly important, and the best role for regulation may be to outline considerations to take into account during development, and offer an analytical framework to identify bias, particularly to ensure accountability and encourage iteration and evolution.

11. One concrete distinction can be made where the provider of the service is a sole provider or operates within an open market. For example, with government run services that structurally do not permit alternatives, Scoring-based decisions can be fully access determinative, and could benefit from a higher degree of oversight. Some private sector services create opportunities for new entrants to step in and seize any opportunities left on the table, creating a very different dynamic. Separately, Presentation-based decisions reflect a variety of market positions, but in some industries effectively render similar results as a result of “filter bubble” problems, which have led to significant discussion but few effective solutions.

12. At minimum, regulation and oversight should consider:

    a. Whether the actions performed or decisions made are simply automation of processes that used to be performed by humans. Operational algorithms will almost always fall into this category.

    b. Whether the provider is a sole provider of a service, like government benefits, or whether there is a robust market with many different algorithms used to make decisions.

    c. The kinds of social and individual impact the decision may have; however, often this will not be immediately clear. This will almost always be a balance.

The scope for algorithmic decision-making to eliminate, introduce or amplify biases or discrimination, and how any such bias can be detected and overcome;

13. The elimination of unwanted bias in algorithms is complicated by the fact that some desired inputs can also be a proxy for other, undesired inputs - highlighting the importance of pairing discussion of decision-making with the discussion of data sets used in that decision-making. For example, one might think that a person's address would be a reasonable input for a car insurance premium algorithm - how likely your car is to get stolen or damaged is affected by where you live. However, address can also be a proxy for ethnic group, because particular groups tend to live together, thus disparately impacting one group. So should this input be acceptable for consideration? The answer is likely context-dependent.

14. While much focus is (rightly) given to potential for discrimination and bias in algorithms, data and algorithms can also find and reverse systematic social bias when used thoughtfully. In fact, in some circumstances, that is a primary potential attractive factor. A well-designed algorithm and data set could isolate considerations of race out of parole assessments, as it has been shown to do in traditionally segregated areas in real estate. And the well-accepted disparate impact analysis is another example of using data and analysis to identify and minimize bias.

Whether and how algorithmic decision-making can be conducted in a 'transparent' or 'accountable' way, and the scope for decisions made by an algorithm to be fully understood and challenged;

15. We believe that "accountability" rather than "transparency" is the best frame in which to consider decision-making. While asking to "show me how it's done" is an appealing idea in principle, in practice it is often very difficult, and not very useful, for the Scoring and Presentation type algorithms in use today. Accountability, on the other hand, is more feasible, and can be evaluated at least in part through known techniques such as statistical analysis.

16. We would encourage the committee to focus on the general goal of ensuring just outcomes, rather than a particular mechanism for doing so. Many decision-making processes are sufficiently complex - and provide social and individual benefit because of

the complexity - that no individual is likely to be able to parse it. Additionally, these algorithms are typically in constant evolution, with ongoing tweaks tested constantly.

17. For some classes of algorithm, data transparency - where a list of test inputs and outputs can be audited - can help to avoid bias. For example, one can check that “ethnic group” is not an input factor in bank loan decisions, and that decisions are not biassed in some way based on the outputs from the data set. There is already a well-established statistical field of “disparate impact analysis” that can analyze outputs by collecting information about inputs and examining both for bias. Indeed, the United Kingdom has already adopted some of these principles in existing anti-discrimination law.

18. The idea of an algorithmic decision or set of decisions being challenged through a legal system makes more sense in some situations than others; in some cases it’s not even clear who would be in a position to bring such a challenge. The algorithm which determines which and in what order items appear in a social media feed can have a significant effect on a person’s mood, and perhaps longer-lasting effects too; but it’s not clear this would provide a basis for anyone to effectively challenge the process or the result.

19. It makes no sense to discuss “algorithms” in isolation, as this misses an essential part of the picture. Algorithms and data are inextricably intertwined, such that discussing one without the other is impossible. Many algorithms are effectively built out of data - the algorithm is simply a training data set, transformed into a form (such as a neural net) which makes it possible to see how well other data matches it. So while it may seem plausible on the surface to make a request for “algorithm transparency” when the algorithm is “add variable X to variable Y, and multiply by variable Z”, it makes far less sense when true understanding requires the release of a potentially huge data set which involves all sorts of data which is subject to copyright, personally identifying, or commercially confidential, or even all three of these at once. If the training is done incrementally, the data set may never even exist in one place at the same time, and so could not be assembled to be released.

20. So there is a risk that well-meaning but misguided rules about algorithmic transparency will place companies in an impossible position by requiring the revelation of data sets which they are legally bound from, or even technically incapable of, releasing.

Methods for providing regulatory oversight of algorithmic decision-making, such as the rights described in the GDPR.

21. The GDPR, specifically mentioned by the committee, does contain certain rights which might well be interpreted to apply to algorithms and their use by companies during data processing. However, the GDPR has not yet come into effect, and open questions regarding the proper application of these rights remains unresolved. We therefore find ourselves unable to fairly assess whether they are an appropriate basis on which to base regulation in this area.

22. As with all technology regulation, we would urge the committee to make any recommendations technologically neutral, forward looking, and designed as far as possible to adapt to the changing environment. There are good reasons why law moves

slower than code, but it does mean that if opportunities for innovation are to be preserved, legal over-prescriptiveness is a danger to be avoided.

Heather West, Senior Policy Manager, Mozilla

Gervase Markham, Policy Engineer

October 2017