

Automation, Correlation and Causation: Launching a Policy Discussion



Following a special workshop convened by the Media Policy Project on 'Automation, Prediction and Digital Inequalities', Tal Zarsky, Professor of Law at the University of Haifa, discusses some of the regulatory and policy implications that arise from companies' use of personal data, with particular reference to wearable technologies.

Technological developments, as well as the rise of new business environments and practices, have brought the analysis of "Big Data" to the forefront of debates about legal and ethical **use of personal data**.

Consider the growing trend of wearable technology and other gadgets (including smartphone apps) that track, collect and transmit health-related data. These tools monitor biometric (heartbeat or temperature) or behavioural (number of steps taken during a set time) factors, thus generating big data. Note, however, that when companies collect and store this information, they often simply aggregate user data rather than identify individual users, although this remains a distinct possibility which generates unique concerns.

Indeed, data from wearables might have high value to a variety of entities which seek to analyse health-related information in order to discover correlations that predict heightened (or lower) disease risk or other behavioural outcomes. **Recent news reports indicate** that employers view the data derived from these wearable sources with great interest, and journalists speculate that insurance companies in the not-too-distant future will use such data to help them set individualised premiums. In these cases, individuals might be harmed and a new policy-based analysis of this scenario is now required.

Companies interested in segmenting and targeting populations on the basis of Big Data often do so by merely relying on a correlation. With such a correlation in hand, companies will strive to predict future instances of heightened risk, or various forms of behaviour in other users. For instance, in some cases, companies might find a correlation between sleep patterns and workplace outputs (a hypothesis recently addressed in the **press**). Thereafter, the company will seek out specific sleep patterns in other employees and on that basis make decisions as to their workflow and job prospects. What the company will *not* necessarily do is strive to understand *why* these factors are correlated with specific outcomes nor the nature of the causal relationship between sleep patterns and work products (if one indeed exists).

Yet is the practice of relying upon correlations to generate predictions without seeking causation normatively acceptable? Though predictions based on a mere correlative analysis are less reliable than ones based on causal analysis, arguably we'll usually be better off letting companies who apply such correlations in practices fail on their own terms, rather than mandating they bake causal analysis into predictive technologies. However, some of the instances involving health data and wearables might indeed call for some form of regulation. Let me unpack this.

A reliance on using just correlations as the basis for action, recommendations to users, or even to make distinctions between them (i.e. discriminating between different individuals in the terms or prices they receive and the coverage health insurance providers offer) generates a variety of concerns. Correlation is often considered to be only the first step of a scientific inquiry and should be supplemented by findings of causation. Indeed, correlations between data factors serve as a tool to formulate hypotheses which are later tested using various scientific tools, such as field or lab experiments.

However, several **commentators** are currently arguing that correlation need not be only the first step, but it could be the last one as well. Relying solely on correlation surely provides several **clear**

benefits. The first is speed. Seeking out causation is time-consuming. Finding correlations can be carried out quickly, automatically and on a grand scale. A closely related factor pertains to costs. Low costs enable startup firms (i.e. companies usually without much in the way of capital or available resources) to engage in data analyses, reveal interesting correlations and offer to act upon them. Engaging in correlation-finding is relatively cheap, as opposed to the other, costly, practices needed to reveal causation. As noted, these can involve engaging in experiments which involve manipulating different factors in the lab. Or, they might require a scientific inquiry in the realms of biology or psychology as to the specific mechanism causing one factor to lead to the other.

Yet these arguments are countered by substantial concerns. Correlation does not show the direction of influence, only that there is a statistically significant relationship between two or more factors. Seeking out causation thus serves as a quality check, allowing analysts to identify faulty data and spurious correlations in data analyses when causality is not apparent. Furthermore, revealing causation can serve additional objectives beyond providing a quality check. Without a theory to understand the prediction, the lesson learned from the correlation cannot be properly *generalised*; it can only be used in the specific context and population where it was found. Developing a theory prior to making data-driven predictions serves as a gauge, making it possible to understand when such generalisations are acceptable and when they go too far.

In addition, relying on mere correlations fails to generate some of the positive externalities and effects derived from a process premised on causation. By searching for causation, we are able to generate substantial knowledge about nature and the human condition. And, when causal rules are made public, individuals who are subject to these predictions will have the ability to engage in self-improvement, at least when the causal factors indicated are mutable. For instance, any correlation (backed by causation) linking specific sleeping patterns to poor job performance might allow individuals to seek out help to address their sleeping problems which could lead to higher work performance. Therefore, while some arguments for relying on mere correlations exist, the case for causation is strong. Balancing these options for data analysis and their normative implications however requires context-specific analysis and examination, which calls for different trade-offs at different junctures.

Generally, we need not aggressively advocate legal intervention and policies which mandate causal inquiries by private companies. Although data analyses which rely only on correlations might be severely flawed, there are also many other business practices which are destined to fail because of faulty management and unacceptable risks. If businesses and their managers choose to undertake a risky and faulty practice, that is indeed their problem. We might instead want to encourage startups to challenge existing scientific paradigms and to introduce new ways of thinking by working with phenomena they cannot yet explain. In addition, it might be unwise to have the law potentially meddling in science more than is needed, as government intervention could be directed by a political agenda or not based on the most up to date research.

Nonetheless, the law might be required to play a role in this debate in the future and even to mandate a causal inquiry in specific contexts, for instance those involving health-related data and any findings which are generated by monopolies or entities assuming a governmental role. It is important for future research to examine and identify the interests of various parties, who society might want to protect and, when needed, for society to apply policy tools to assure that Big Data analysis goes beyond mere correlations.

To conclude, consider the following examples: in health-related contexts causation rules might be required given the importance of such knowledge to science, as well as the risks of errors in the conclusions that follow, if generalisation is carried out wrongfully. However, in situations related to credit allocation, in which correlation-driven business models unfold in a competitive setting, the risks of mere correlation might be limited, and therefore regulatory intervention might not be called for.



This blog gives the views of the author and does not represent the position of the LSE Media Policy Project blog, nor of the London School of Economics and Political Science.

This post was published to coincide with a workshop held in April 2016 by the Media Policy Project, 'Automation, Prediction and Digital Inequalities'. This was the third of a series of workshops organised throughout 2015 and 2016 by the Media Policy Project as part of a grant from the LSE's Higher Education Innovation Fund (HEIF5). To read a summary of the workshop, please click [here](#).

June 10th, 2016 | [Digital Inequalities](#), [Featured](#) | [0 Comments](#)

☺

