Transcript

CHATHAM HOUSE
The Royal Institute of International Affairs

The Hidden World of Algorithms

Luke Dormehl

Journalist; Author, The Formula: How Algorithms Solve All Our Problems, and Create More

Dr Juan Pablo Pardo-Guerra

Assistant Professor, LSE

Dr Daniel Neyland

Professor, Goldsmiths University London

Chair: Zoe Kleinman

Technology Reporter, BBC

21 April 2015

10 St James's Square, London SW1Y 4LE T +44 (0)20 7957 5700 F +44 (0)20 7957 5710
www.chathamhouse.org
Patron: Her Majesty The Queen Chairman: Stuart Popham QC Director: Dr Robin Niblett Charity Registration Number: 208223

Zoe Kleinman

Good evening, ladies and gentlemen, and welcome to our panel discussion 'The Hidden World of Algorithms'. My name is Zoe Kleinman. I'm a technology reporter at the BBC and I'm chairing tonight's event. Before we get started, you'll be delighted to know that it is on the record. If you are tweeting, please use the hashtag #CHevents. But please keep your phone on silent, obviously. If you can, please join us for a drink at the reception afterwards. It should be at around 7:30.

We're calling this event 'The Hidden world of Algorithms', but actually I'm not really sure that they are so hidden. These little 'if this, then that' programmes pop up in almost every aspect of modern life as we're going to hear tonight. Just today, I was on TV talking about Google updating its algorithms in order to prioritize mobile friendly websites in its search rankings.

Just two years ago, whenever I did anything about algorithms on air, I'd get a note from the producer which says something like, 'Zoe, if you've got to mention algorithms, for God's sake, at least tell us what it is.' Then I took a few months off last year to have a baby, and guess what? Apparently, in my absence, the world and perhaps more crucially the media, found out for itself what algorithms are because I no longer have to elaborate.

But if I had to pick one vital ingredient that our algorithms still lack, I think I'd call it tact. I won't repeat what I said when a supermarket sent me some vouchers recently based around my shopping habits. The first offered 25 per cent off a bottle of champagne. Great, I thought. The second, extra loyalty card points – whatever they are – on a box of chocolates. Okay, I can have that. But the third was a free test for type 2 diabetes. Logically, I know I can't blame the algorithm for that. But nonetheless, I can tell you that all three went in the bin.

Without further ado, let me introduce to you tonight's speakers. First up we're going to hear from Luke Dormehl. He's a journalist, filmmaker and author of the book, *How Algorithms Solve All Our Problems, and Create More*. Well, I think I've just given a good example of that. And also, another book called *The Apple Revolution*. He's had films screened on Channel 4 – the competition, and at the Cannes Film Festival, and he's written for *Wired* and *Fast Company* among other publications. So Luke is going to talk about algorithms and their impact on the modern world.

Next, we're going to hear from Dr Juan Pablo Pardo-Guerra, who's an assistant professor in Sociology at the London School of Economics and Political Science. He's going to focus on the impact of algorithms in the financial sector. His expertise lies in the links between markets, politics and technology, and he's published research in several journals, including the *Journal of Cultural Economy* and *Technology in Society*.

Finally, Dr Daniel Neyland to my far left, is professor of Sociology at Goldsmiths, the University of London. Dr Neyland's broad range of research interests, and this fascinated me, include traffic management and malaria. I'm not quite sure what the link is. He's also an expert on governance and accountability. You'll find his research in publications like *Organizational Ethnography* and *Privacy, Surveillance and Public Trust*. He's going to consider the governance and privacy issues surrounding algorithm use.

Luke Dormehl, if you'd like to start, please.

Luke Dormehl

Fantastic. Thank you very much all of you for coming. I'd like to start with a brief story. Back in the 1980s, William Goldman, the Hollywood screenwriter, was working on his autobiography. Goldman was a very successful screenwriter. He wrote *All the President's Men* and *Butch Cassidy and the Sundance Kid*, among other films.

However, when he was asked to comb through his years in the film industry for something profound to say or something that he could share with readers, all he could come up with was the idea that when it comes to movie making, nobody knows anything. 'Why,' he asked, 'did every studio except for one turn down *Raiders of the Lost Ark*? Why did the biggest studio of its time turn down *Star Wars*?' Because, according to Goldman, nobody knows anything.

Now, this isn't mean to demean people who work in the entertainment industry or to accuse them of missing enormous opportunities, although of course in this instance they did. Anyone who works in entertainment I daresay has his or her war stories about the hit that got away or the flop that seemed like it would be a hit. There are sure-fire winners which become losers and apparent losers which turn around and become winners.

There are niche films which appeal to everyone, and populist films that appeal to no one. Hardly anyone has an unblemished track record, and trying to make sense of it all is really baffling to even those humans who are paid large sums to do exactly that. There is a company, and it's actually not located too far from here in London, which claims to be able to help.

This is an algorithmic script consultancy called Epagogix, which uses algorithms and machine learning tools to tell some of the biggest studios in Hollywood which films will earn them the most money at the box office. More impressively – or scarily, depending on your perspective – Epagogix can even make creative decisions by singling out places in the script where the potential yield isn't what it could be and then recommending that screenwriters work on those areas.

This is essentially the same kind of data mining that you've probably seen, or that we have seen in multiple industries. If you've watched the movie, *Moneyball*, for example, then you've seen how it can be used to turn a baseball team of apparent losers into winners. But there are plenty of other areas where algorithms are also used.

Before I go any further, I should really define algorithms. An algorithm, for those unfamiliar with it, is really at its highest level a series of step-by-step instructions designed to complete a certain task in a finite amount of time. You could, for example, think of cooking a recipe as an algorithm, since we start out with a set of ingredients, we follow a series of guidelines, we know that when a certain colour or temperature is achieved, that we've achieved our end result.

An algorithm, in other words, is a method of turning inputs into outputs. But if their description is straightforward – and of course there is a bit more complexity than I've described – their application is certainly anything but straightforward.

Today it's difficult to think of a facet of human life that hasn't been transformed by algorithms. They sort, filter and select the information that's shown to us on a daily basis. They're responsible for the search results that Google shows us, the news stories that Facebook thinks we'll be interested in, the partners that dating websites think we'll be compatible with, and the goods that Amazon thinks we'll want to buy. They're changing all the time.

Last year, Amazon for example, was granted a patent for something called 'anticipatory package shipping', which means that it could conceivably send out products that customers are most likely to purchase, before a customer actually orders them.

They're applied in law enforcement, too. You have predictive policing, for instance, which can not only forecast the likelihood of a certain crime taking place, but also who's likely to commit it, what type of crime, when, where they're likely to commit it. Yesterday I came back from a weekend away. When I boarded the plane, I was quite possibly subject to algorithmic profiling, deciding whether I was likely to be a terror suspect, using the same analysis of seemingly unpredictable data that Epagogix might use for films.

Of course, any force which governs our lives – especially one that comes about as undemocratically as the rise of the algorithm, in the sense that these are decisions made about what is relevant or desirable or interesting made by a relatively small number of software engineers – any force like this carries a certain number of problems or challenges.

One of these is the idea that these formulas or algorithms really aren't discoverable by us to an extent. We, as you sort of touched on in your opening comment, we're kind of aware of this idea of what an algorithm is perhaps, but for the most part, a lot of the algorithms we rely on are black boxed, meaning that we really see the inputs, we see the outputs, but that sort of messy middle stage of the sausage making process is invisible to us. It's obvious why, when you're talking about billions of dollars being at stake for Amazon and Apple and Google.

In other instances, there may be security reasons why you might not, why it may not be advisable to release the source code for one of these algorithms. This means that an individual targeted by an algorithm as a potential terror suspect, for example, might find themselves questioned for hours or forced to miss flights without ever finding why they were targeted to begin with.

What we rely on in these situations is the idea that algorithms strip away the subjectivity which leads to prejudice. An algorithm like the one that's used for box office numbers isn't going to suggest the hiring of Will Smith because it wants to hang out with him at a Hollywood party. Nor should an algorithm used in law enforcement let a person off a drink driving charge because they happen to be a powerful politician.

'Our users trust our objectivity,' says Google, summing up its attitude in a manifesto it calls '10 things we know to be true'. Science fiction author Arthur C. Clarke once suggested that any sufficiently advanced technology is indistinguishable from magic. Just like photography appeared to people a century ago to be a medium unaffected by human tampering, so today do we look at algorithms as the embodiment of truth.

On some level, most of us like the idea that we can enter inputs, expect outputs, and not worry about what happens in between too much, believing that an algorithm can do something as seemingly straightforward as tell us which stories are relevant on our Facebook feed. Relevant to whom? It speaks volumes about our desire for easy answers.

There are of course plenty of examples of algorithms gone wrong and I'm sure we'll touch on a few of those this evening. I'm particularly interested, I suppose, when inevitably we see our first death that's the result of a self driving car. I think we're going to start discussing this perhaps a lot more. But more crucially than what happens when an algorithm goes wrong, is really the question of what happens when it works correctly.

Programmes to personalize our experiences online, algorithms flatter our personal mythologies and sometimes confirm our existing biases. The famous example, of course, is the liberal who types BP into their Google browser is much more likely to get information about the oil spill in the Gulf of Mexico than the conservative, who's more likely to receive investment information.

Beyond this, lurk questions about techno replacement and the idea that within our lifetime, many of our jobs are essentially going to be handed over to automation. This might be a company like LegalZoom, for example, which promises to carry out a lot of the work that a human lawyer might currently do for a fraction of the price.

As with Epagogix, similar ideas about algorithmization may be one day applied to the creative process, writing scripts or generating paintings or novels that are calculated to please us. House of Cards, the popular Netflix television show, was assembled by using algorithms to analyse which TV programmes a desired audience most enjoyed. The creators concluded that their desired audience enjoyed BBC dramas, Kevin Spacey and the director of *The Social Network*. So they got David Fincher to direct Kevin Spacey in a remake of a very successful BBC political thriller. It didn't take too long for the awards to start racking up.

Ultimately, I don't have all the answers to these questions. I don't think any of us do. I like to quote the technology theorist, Paul Virilio, on the subject of technology. He said that the inventor of the ship was also the inventor of the shipwreck. To that end, we might add that technology is not good nor bad, but neither is it neutral. At the end of the day, it's a matter of appreciating that if algorithms claim to provide all the answers, we have to start asking the right questions. Thank you very much.

Zoe Kleinman

Thank you, Luke. Next we're going to hear from Dr Juan Pablo Pardo-Guerra, who's going to talk to us about the financial sector and algorithms.

Juan Pablo Pardo-Guerra

Thank you very much. I want you to look at the clock on the wall for a little bit. One, two, three, four, five. In five seconds, the most sophisticated trading systems in the financial world today can execute up to 50,000 individual trades, give or take. This is executed by each of these systems, and there are many of these systems around. In five seconds, you can potentially have one million trades in global financial markets executed automatically.

To put this in context, I want you to think of human cognition. While human cognition exists in the threshold of 200 to 400 milliseconds, so thousandths of a second, trading systems and finance today operate in the realm of 100 microseconds, so millionths of a second. This is far beneath any of the processing capabilities of both human minds or human bodies. Finance today is something that moves much faster than our capacity to understand the world that understands us.

This is what is often referred to as high frequency trading, an activity that accounts for roughly 50 to 60 per cent of the trades executed in most European and American financial markets or stock markets primarily. It is as we all know, a topic that has been at the centre of public debates on the standing of financial markets in the recent past and that is highly controversial. If you want to see levels of controversy, just ask Michael Lewis about how the industry that is involved in high frequency trading received the book. It's highly controversial.

What is the origin of high frequency trading? Well, there are three distinct cores or historical trajectories that are associated to the emergence of this type of activity. The first is the wide use of sophisticated communication technologies in financial markets. In the past, markets were relatively isolated and disconnected pockets of trading. So you had stock exchanges connected by runners, telegraphs and phones. That was the golden day of finance according to some people in the industry.

Today, markets and trading systems are linked through high speed networks through which petabytes of information travel every year seamlessly through regulatory borders and through different organizations of numerous sorts. That's the first big change, information or communication technologies, rather.

The second was a change in the type of people working within financial institutions. So today, you find physicists and statisticians employed in your average financial institution, whereas in the past, you would find more gentlemanly brokers and jobbers operating within those firms, in a sense. Those were the individuals that populated the floors of finance in the past. So that's the second change, a change in the expertise within the market.

The third, of course, is the possibility of delegating simple tasks, like looking at the average price over the last five milliseconds, to a computer that through an algorithm then decides on the basis on some vast statistical analysis, what to trade, when to trade it, and where. Every single one of those decisions is delegated onto an algorithm. Of course, this is cheaper than having someone sitting in front of a phone and it's much more efficient. That's one of the reasons why this was taken up.

Of course, this is also where much of the attention has gone to in the recent past, the recent years, to the automation of financial markets through algorithms. The question really is, should we be worried about the expansion of algorithms or the dominance of algorithms in the world? That markets are now made by algorithms rather than by human beings?

I'll be slightly controversial for a sociologist, and I'll say that actually it's not that big a deal. It's actually quite uncontroversial in many ways.

Siding with the defenders of high frequency trading, most economists and sociologists who have studied these transformations agree and suggest that automation has produced what are perhaps the cheapest, most easily accessible and most efficient stock markets in human history. So we have very good stock markets that work relatively well, that behave fine and that's generally good. This has to do with the process of automation.

Studies have also shown that algorithms, unlike what some argue, do not destabilize the market or introduce spurious volatility into prices, but actually allow information to be reflected better in prices, because they are so fast and they are so efficient. That's one of the reasons. They work, and they work well.

Indeed, to say that market algorithms are a problem would be to fail to recognize the history of markets and the history of decision-making within markets over a longer period of time. So algorithms are very old inventions. They've existed for quite some time and at a very basic level, they're just instructions. So like recipes in a cookbook.

In markets, they've been around for at least two centuries. If you go back to the 18th and 19th century, you'll find pamphlets and books that have recipes for financial success. They would tell investors how to trade, what to do in the market, when the market was moving in particular ways.

In the early 20th century, you had chartists or technical traders that, on the basis of some movements in the market, both past and present, would give you instructions on how to trade that were essentially algorithmic. These are algorithms, rules or instructions that have existed for at least a century in the market.

So perhaps the anxiety with algorithms in finance has to do with the idea that they act on their own. But what we have to understand here is that algorithms aren't actually independent. At the end of the day, their operation is tied and grounded upon decisions that are taken within financial organizations every single day, hour after hour.

For instance, in high frequency trading firms, algorithms have to be adjusted constantly. They don't just roll them out and let them live happily in the market. They have to tweak them. They have to revise them. They have to tune them in order to take into account the difference states of the market that are decided, to some extent, on the judgement and experience of traders sitting in front of screens. So there's a human at the end of every algorithm, to some extent, and a complex organization.

This is something we see also in things such as Google's page rank algorithm, which isn't just any single algorithm, but it's a series of organizational decisions on what the system

should so in terms of the changing structure of the world wide web. It's an algorithm. It's a series of instructions that are constantly revised. We should perhaps remember what philosopher Ludwig Wittgenstein suggested, which is this idea that rules are never sort of unambiguous. They are always in the process of being reinvented and reapplied and reinterpreted, in a sense.

So what is it about algorithms that is actually important from perhaps a sociological perspective? Well, one of them is that they are coupled to this change of expertise in the market, and that is something that hasn't been mirrored by regulators. The issue here is that regulators in financial markets tend to come from professions – lawyers and economists – that do not have the technical skills perhaps to deal with the changing nature of market infrastructures.

Perhaps the rise of algorithms calls for a rise of a different type of regulation, one that is closer to telecoms regulation than to regulation of financial markets as is traditionally exercised, in a sense. This calls for thinking about the role of data in modern societies, and all the complexities that are associated to that. That's one of the things that algorithms present to us – the need to change perhaps regulators. You would never take your broken computer to a lawyer so that it's fixed by him or her.

The second is of course that the rise of algorithms may actually be quite disruptive to the industry and a small example happened a week ago. So last week, a new company hit the market in the New York – Virtu. This is the first high frequency trading company that went public, that issued shares to the public in New York.

They were capitalized according to the market at 2.6 billion, which is minute by financial standards, and they are minute. They have about 150 employees worldwide, yet despite that small size, despite the fact that they're a small challenger to the incumbents, they have higher volumes of trading than most conventional investment banks.

This small company may actually change the business models in finance or at least in one important sector in finance. It might induce a new world in banks where banks will be smaller, more efficient, and who knows, perhaps a little more trustworthy in a sense. That is one of the things that might emerge out of this change induced by algorithms.

All in all, it's not a bad picture. It's actually a very positive one if you think of it from this perspective. Thank you.

Zoe Kleinman

Finally, let's here from Dr Daniel Neyland. He's going to talk about governance and privacy issues.

Daniel Neyland

Thanks for inviting me to take part in this debate. In order to get some discussion going, I'm going to argue against Juan Pablo. So let's see how we get on.

I guess your two kind of uncontroversial points about algorithms that we could start thinking about is that algorithms have been around for a long time, as Juan Pablo said. You can trace the term back to about the ninth century, and for decades there has been discussion of algorithms in relation to forms of machine data. So maybe there's nothing new in that sense.

Algorithms on their own, also maybe don't necessarily do very much. As we've heard, they're a kind of set of 'if, then' rules. They're a set of conditions and consequences. Sometimes on a piece of paper, sometimes they're on a computer screen. On their own, they don't necessarily produce much of a consequence.

Why the concern? Why this talk of the hidden world of algorithms? Why a kind of rise in questions about what the nature of an algorithm is? Why should we be concerned about new forms of regulation?

I'm a sociologist who carries out research in the field of crime control, broadly conceived. So recently I've been working with people like airport operators and train station operators and large technology firms interested in the use of algorithms for sifting through things like digital video data from CCTV systems to pick out particular kinds of activity from that data, and then make decisions based on that.

Things like people moving in odd ways, abandoned luggage, large crowds. The idea is you pick out any kind of reasonable activity that you might want to pick out as relevant, you look at the way the stream of digital video data would reflect that change in activity and then you design an algorithm and an associated set of software code that would then allow you to pick out that kind of behaviour.

But the algorithms themselves are just sets of 'if, then' rules. If a certain pattern of behaviour is uncovered, then a response is required, for example. The algorithms do nothing on their own. They only work as part of existing video surveillance set-ups. They need an infrastructure. They need cameras. They need people. They need technology. They need a space like the airport or train station that's reasonably defined for this to happen.

But perhaps this is why we should be concerned. Perhaps this is why we should be concerned. This is what is new, I think. The existing things like video surveillance set-ups are being turned into algorithmic systems, combinations of people, technology, databases and so on linked together with algorithms and associated code.

I'm going to ask four questions that I think help us to get to the heart of maybe what some of the concerns are or some of the things that we might need around this move to develop algorithmic surveillance systems. First question: are we happy to delegate to algorithms and algorithmic systems the role of deciding what is and is not relevant, what might require a response in somewhere like an airport or a train station?

Well, maybe. There's an argument in favour of this. Surveillance systems are getting larger, monitoring sort of 200 to 500 cameras typically. It's quite a tough job for a surveillance operator to do. It's difficult to look at that many things at once. So perhaps algorithmic technologies that can sift through digital data provide assistance, that steer

operators towards particular activities, might be helpful. From the perspective of the operators of these systems – airports, train station operators – maybe these technologies are valuable because they might be able to employ less staff. Fewer staff, lower costs to cover more cameras.

But maybe not. Maybe it isn't such a good thing. It seems quite difficult to hold an algorithmic surveillance system to account, to understand precisely the nature of the algorithms that are involved, what they do, with what consequences. Some of these algorithms seem to inadvertently target particular groups more than others, particular individuals more than others, particular types of styles of behaviour more than others. So a technical prejudice seems to seep in at certain points. Maybe we should be concerned about that.

The second question: are we happy to allow algorithms to create ever more associations between data? To inform decisions whether or not to intervene. For example, to link things like streams of digital video data from CCTV systems with things like passenger record databases, criminal record databases, other databases that could be made available to the system. Maybe. Maybe we should be happy about that if we're particularly worried about security, crime, safety, acts of terror, and we perceive that those kinds of combinations between datasets might help, maybe we should be happy about that.

Also maybe not. If we're concerned that this approach assumes that criminality is the kind of default position, that everyone should be assumed to be a criminal until proven otherwise, their quality of life might be inhibited, our freedoms infringed, consequences uncertain, and no particular clarity on where this might end. Maybe we should be concerned.

Just how far should we go in creating these associations between different forms of data? How much data should be collected about us? Should we have concerns about things like government track record on holding large sets of data, losing it, making incorrect and inappropriate decisions based on that data? Maybe we should be concerned about that.

Third question: are we happy for algorithmic systems to sit within current legal, governance and accountability structures? Well, maybe. I think there are some positive arguments to say here. In some ways at the moment, it's still quite appropriate to look at the decisions made by CCTV operators, police officers and so on. Maintaining that kind of human centred rather than algorithmic centred form of regulation is still quite important. There are still often people making decisions at the centre of these systems.

Also, algorithmic systems are subject to data protection regulation in the same way that other sorts of data regimes are and so those sorts of rules are still important. Rules around things like the necessity and proportionality of data collection still seem appropriate for algorithmic systems. So some of the current regulatory landscape is still relevant.

At the same time, the current European Data Protection Act is from 1995, although of course there are discussions going on to replace it at the moment. It still seems kind of stalled and held up with various discussions going on. Since 1995, it seems to me that the nature of data has changed quite significantly, in terms of the scope and scale of data, in

terms of the way data can be read, mined, scraped, associations created between different forms of data, the types of consequences that can flow from data. They all seem to have changed quite significantly.

Perhaps we need some updating of those kinds of regulatory principles, but also the kinds of mechanisms for how we might govern and hold to account algorithmic systems. So perhaps, as Juan Pablo was talking about, we need to think about new forms of expertise. We need to somehow get in at the design stage in shaping the kinds of algorithmic systems produced. Perhaps we need to insist on things like still having humans as decision-makers involved in these systems. Perhaps we need to develop a capacity to understand and track some of the consequences of algorithmic rules and interventions.

The kind of research that I do and other people do can play a role here. But research is kind of small scale and often fixed term, short term. For every one system that I might look at up close, there are probably 100 other systems being developed. To some extent, research can't really play the role of doing governance and regulation here.

So how might we govern algorithmic systems in an ongoing way? One possibility that I'm not actually particularly keen on, but one possibility is to expand the capacity of algorithmic systems in order that they might participate in their own governance and accountability. Yes. So I'm kind of torn about this.

One possibility is to design algorithmic systems in order that for example they automatically delete data after a set time, so that it cuts down on the amount of data that gets stored, automatically deletes the data that doesn't get looked at. Perhaps puts in place software and code that prevents particular associations being made between different databases. Or you can create algorithmic systems that report on their own activities.

For example, you can introduce things like access management systems, digital CCTV systems so that people have to log in and perhaps even have to give a justification for why they're looking at data. That can become an automated log that goes to, for example, data protection officers so they can see how a system is being used and how decisions are being made. That's maybe one option. Perhaps there are other options that the audience might like to suggest.

The final fourth question: are we happy that algorithmic systems work well enough and are fulfilling their aims to be given the responsibility to shape interventions in crime control and security situations? I would say no. I don't think they work well enough. From the research I've been doing, it seems that there are probably too many false positives from a privacy perspective, too many people and things that are identified when they shouldn't be, and too many false negatives from a security perspective, too many things that ought to be looked at which are not by these systems. I'll stop there.