

The Ethics of Data Science Conference
Centre for Translational Data Science, University of Sydney
28 March 2019

Transcript of Oral Remarks & Accompanying Slides
Dominique Hogan-Doran SC



Computer Says 'No' -
Automation,
Algorithms, &
Artificial Intelligence
in Government
Decision-Making

Dominique Hogan-Doran SC

The Ethics of Data Science Conference
University of Sydney, 28 March 2019

Liability limited pursuant to a scheme approved under professional standards legislation



Listen on Soundcloud:

<https://soundcloud.com/user-176287914/the-ethics-of-data-science-conference>

So why am I here? Why is a lawyer here? People often ask that question. So, it comes in this context.

This is a presentation of a paper that I was asked to prepare for the Judicial Commission of New South Wales in 2017, which was published in the *Judicial Review*, which is this text here.

[moves to next slide]



- ❖ Many administrative decisions that used to be based on human reflection are made automatically.
- ❖ Bureaucracy is no longer what it once was.
- ❖ How will public law's mandates of transparency, fairness, and accuracy be guaranteed?
- ❖ Increasing pervasiveness of technology assisted decision-making makes the public law challenge particularly acute.
- ❖ The full basis for algorithmic & robotic AI decisions is rarely available to affected individuals
- ❖ We require more governmental transparency, clear & effective regulation, & widespread awareness of dangers & mistakes already occurring.

3

I was asked to prepare and to present it, in the context of doing what you're doing here, which is to create a conversation, and to introduce these ideas and these technological developments to a group of people – in this case, who are technologically illiterate - Judges.

And so, what I had sought to do in this paper is to introduce to them the existence and the growth of automated decisions & technologically assisted decision-making in the government sector. To help them understand how it was developing, and the kinds of ways it would come to them for consideration.

Now this was a one-hour presentation, and set of slides that I have attempted to edit down. Lawyers tend to talk a lot. We also tend to write a lot. So there will be a few slides that I will skip straight through, but it's all on my website if you want to look at them further. *[moves to next slide]*

Why use technology to assist government decision making?



- Modern bureaucracy is an ideal candidate for algorithm-based, automated habitus:
"Bureaucracy is an organisational structure that is characterised by many rules, standardised processes, procedures and requirements, number of desks, meticulous division of labour and responsibility, clear hierarchies and professional, almost impersonal interactions between employees" - Max Weber, *The Vocation Lectures*
- Optimize bureaucratic size and depth
- Process large amounts of data:
 - ✓ more quickly
 - ✓ more reliably
 - ✓ less expensively
- Promote [but not guarantee] lawful decisions:
 - ✓ ensure decision-makers act within limits of powers
 - ✓ mitigate risk of irrelevant considerations
 - ✓ avoid bad faith
 - ✓ avoid improper motivations

4

Explaining to judges and explaining to people like you about why technology would be used [in government decision making], is not surprising at all. Because it can assist in processing large amounts of data quickly and reliably and inexpensively. And also, it is held up as being an advantage to promoting lawful decisions.

[moves to next slide]

When best to use technology to assist decision making?



YES

- ✓ Make a decision
- ✓ Recommend a decision to a decision-maker
- ✓ Guide user through relevant facts, legislation and policy and/or close off irrelevant paths
- ✓ Provide decision-support systems eg link relevant legislation, case law and policy, or useful commentary
- ✓ Self-assessment tool: provide preliminary assessments for individuals, internal decision-makers

No

- Discretionary decision
- Decisions involving discretionary elements
- Where would fetter decision-maker in exercise of power by recommending or guiding decision-maker to a particular outcome

Administrative Review Council, [Report no. 46](#), *Making Automated Assistance in Administrative Decision* (2004)

5

But the traditional view in Government and in the legal sector as to what use can technology play, has been quite traditional and restrictive. That is, it can be used essentially as *assistance* in making a decision, but not making a decision itself.

And the way that that is amplified in what we call 'administrative law', is in the world of discretionary decisions.

The problem is, just taking a step back from it, is that most of the work that was done was done at a time in the early 2000s, - 2004, 2007, - when the only expert systems were ones which were rule-based systems, logic systems. There wasn't what there is now - where there's machine-learning and at least robotic artificial intelligence, which can substitute for a human decision-maker, substitute the exercise of discretion, where a large range of factors can be taken into account, can be part of the consideration. That just wasn't considered, let alone contemplated.

[moves to next slide]

Australian Government Information Management Office (AGIMO)
Automated Assistance in Administrative Decision-Making Better Practice Guide, 2007

- ✓ Assess the suitability of automated systems to deliver improved business outcomes for an agency
- ✓ Establish appropriate project management and governance of automated systems projects
- ✓ Ensure that the design of an automated system has regard to future requirements (such as maintenance and audit) and complies with privacy legislation
- ✓ Ensure the continued accuracy of an automated system (including where there are changes to the underlying legislation, policy or procedure)
- ✓ Ensure the transparency and accountability of the system and its accompanying processes, and
- ✓ Implement and maintain automated systems appropriately

<https://www.oaic.gov.au/images/documents/migrated/migrated/betterpracticeguide.pdf>

6

And so, there's this *Better Practice Guide*, and I'm going to come back to it more, which was issued [by the Commonwealth] in 2007 which is as far as I know it's still not been updated. And the last two points are the key points I want to draw to your attention. What was the message being sent to Australian administrators? You need to:

- ensure the transparency and accountability of this system and its accompanying processes, and
- implement and maintain those automated systems appropriately.

Now that's a straightforward proposition when one can record in a logical step-through process, the way in which computer programming has assisted. But today, answering those two dot points is a far more complex challenge for administration in Australia and throughout the world.

So, in this context then, I want to take you quickly through two slides, one at the high level and one at the granular level. Because what I'm attempting to do today is introduce to those of you who are data scientists, who are each experts in your particular area, some insight into what is the context in which today artificial intelligence ultimately, but at least for the moment machine learning & profiling, all of those technological advances, what's the context it has to play out in?

So if you're creating business rules, or including a technological input into a government decision-making system, you've got to deal with some very difficult complex concepts which will be foreign to you. Just like to me, even though my father is a computer programmer and I've grown up amongst it, to me the technology is difficult at least.

[moves to next slide]



- Substance & breadth of legislation
- Structural complexity of legislation
- Semantic complexity of legislation
- Exercise of discretion (risk of fettering)

7

So at the high level, these are the things that have to be considered.

[moves to next slide]

***Lawful government decision making can be hard:
public law informs us that risk of errors of law abound***

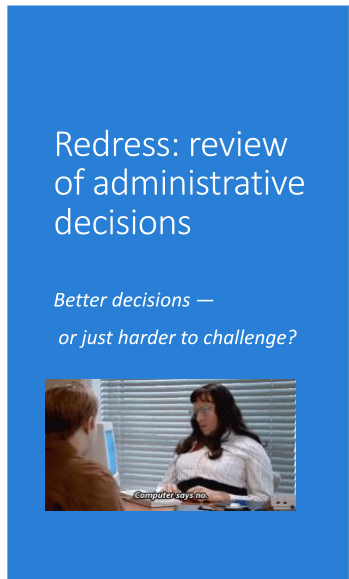
- Identify wrong issue/ask wrong question
- Fail to address the question posed
- Ignore relevant consideration/rely on irrelevant consideration
- Apply wrong principle of law
- Breach mandatory statutory procedure or obligation
- Engage in unfair treatment
- Fail to give requisite notice/opportunity to respond
- Fail to give access to all information/documents relied on
- Mislead as to intention/fail to adhere to statement of intention
- Mistake the evidence, misunderstand the claim
- No/absence of probative evidence
- Based on evidence not meeting applicable standard of proof
- Insufficient evidence due to inadequate inquiries
- Disproportionate/excessive weight to factor
- Use reasoning illogical or irrational
- Lack evident and intelligent justification

8

But at this lower level it's even worse.

So when errors are made in decisions, there's lots of different ways you can make a mistake. Lots. There's a whole world of law that has dealt with this. And it's not like the law is going to change quickly to deal with your problem.

[moves to next slide]



Redress: review of administrative decisions

Better decisions —
or just harder to challenge?

Computer says no

Merits review

- Merits reviewer stands in shoes of original decision-maker. Not bound by original decision
- Merits reviewer not bound to take into account any findings or conclusions offered by computer system or relied upon by human decision-maker
- If same technology used by first-instance decision-maker as well as merits reviewer, is there a risk that a flaw or error in technology will carry from initial decision-maker to merits reviewer?

Judicial review

- Automated decisions capable of being judicially reviewed
- Automated systems must be clear about their reasoning process and what questions / materials were considered
- Court may require evidence from expert witnesses in order to understand code embedded within technology and reasoning process used by automated system

9

I want you to think about the *citizen* versus the *government*. The citizen is unhappy with the decision that's been made. The citizen can ask for that decision to be reviewed.

By bringing technology into the decision-making process, *are we helping better decisions be made? Or are we just making it harder for citizens to challenge them?*

And the point I want to make here is, the last dot point under “merits review”. In the merits review, the decision-maker stands in the shoes of the original decision-maker. Well, how does a human stand in the shoes of a robotic artificial intelligent decision-maker?

And in judicial review, judges want to see: how did the decision get here? What was the process?

So someone like me who's a barrister, who's either asked to defend a government decision, or asked to assist an individual who wants to challenge a government decision, says, "*All right, tell me how this came about?. How do I explain this to the judge, and identify where the error has gone, and where it needs to be corrected?*" So I may need to call an expert witness. That means I might call you. "*Come and explain to this judge here how the hell you got there. What went wrong? Or, Why did it go right?*"

[moves to next slide]



- 2016 launch by Department of Human Services ('DHS') – Centrelink of new Online Compliance Intervention ('OCI') system for raising/recovering social security debts
- OCI system was a complex automated system rolled out on a large scale within a relatively short timeframe
- Non-employment Income Data Matching Project designed to enable DHS to use a data-matching algorithm that measured income on an annual basis and divided it into equal fortnightly instalments
- Matched income data DHS collected from customers with tax return data reported to the ATO by employers
- Differences in income notified to customers by letter, reversing onus to correct/update DHS records.

10

Robodebt - a good case study! Why is it a good case study? Well, if nothing else, it was a great consciousness-raising moment. Because suddenly people felt, "*Oh yeah, computers are involved in government decisions. It's not just some nameless bureaucrat. Now it's a nameless computer.*"

And what was the problem? Essentially it was the problem that it was matching income data collected by the Department on a fortnightly basis, (Centrelink payments are paid fortnightly) and matching it against tax return data which is reported by employers on an annual basis.

[moves to next slide]

Centrelink's "RoboDebt" controversy

Not a problem of automation *per se* – ie adoption of automated systems in which efficiency was a primary goal.

Rather, source of problem was Department's flawed implementation of an algorithm which was subsequently automated, without clear processes & explanations to deal with exceptions.

Senate Inquiry (2017) 20% of debt notice recipients demonstrated no debt was owed

Commonwealth Ombudsman Report (2017):

- critical of fairness & reasonableness of roll-out, poor planning & risk management
- inaccuracies of income averaging algorithm, especially for casual workers
- poor communication with vulnerable customers (misunderstood need to engage & supply information, correct errors)
- fettering of discretion re waiver of penalties (a 10% recovery fee had been embedded in the business rules encoded in the OCI system)

11

And what was the controversy?

It wasn't a problem of automation *per se*. Not just because automated systems had been used to make a more efficient process. That's fine.

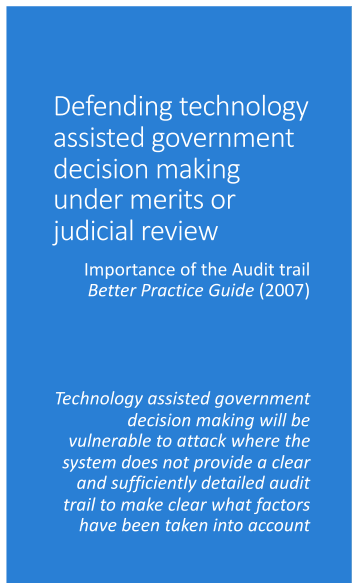
Rather the problem was the implementation of an algorithm which was subsequently automated, and then there were the broader contextual problems. It didn't explain it. It didn't appreciate. This is in its original implementation, there's been major modifications since. It didn't understand its target audience, which had groups of very vulnerable people in it. The Commonwealth Ombudsman's Report was critical of the rollout. Done too quickly, poor planning, poor risk management.

You can understand how the algorithm which was then extrapolated out from a fortnightly basis. Someone might work just a couple of times a year will have disparate incomes because they're working as a casual worker. And then when people receive the letter, their first response was not, "*Oh the computer don't know, the computer was wrong.*" It was, "*I've got a problem.*"

So when you come to defend, when someone like me has to come and defend the technologically assisted decision that's under review, what am I going to be looking for? Well at the moment all that has really been identified is: show me the audit trail. Show me how it got there. All right.

So these are the things I'd be looking for. This page and the next page.

[moves to next 2 slides]



Defending technology assisted government decision making under merits or judicial review

Importance of the Audit trail
Better Practice Guide (2007)

Technology assisted government decision making will be vulnerable to attack where the system does not provide a clear and sufficiently detailed audit trail to make clear what factors have been taken into account

- Does the automated system have the capacity to automatically generate a comprehensive audit trail of the administrative decision-making path?
- Are all the key decision points identifiable in audit trail?
- Are all the key decision points within the automated system's logic linked to the relevant legislation, policy or procedure?
- Are all decisions recorded and accessible by the system's user, a reviewer or an auditor?
- Can audit trail generated by the automated system be easily integrated into a notification of the decision (including a statement of reasons or other notification) where required?
- Is audit trail secure from tampering (to provide protection and data integrity)?

12



- Does audit trail include a comprehensive and printable modification history including:
 - who created the document (with time and date recorded)?
 - who has modified the document (with time and date)?
 - a record of what was modified?
 - for privacy and commercial-in-confidence matters, who has viewed the document (with time and date)?
 - who made the final decision (with time and date?)
- Does audit trail start by identifying the authority or delegated authority identified in legislation?
- Does audit trail show who an authorised decision-maker is?
- Does audit trail enable the recording of human intervention in automated processes, for example recording who is authorised to exercise intervention?

13

But you and I know ... This is one of those slides I'm not asking you to look at for long.

[moves through next slide]



Benefits of profiling	Risks of profiling
Better market segmentation	Infringement of fundamental rights and freedoms
Permits analysis of risks and fraud	Certain sectors of society may be underrepresented — eg older generation/ vulnerable individuals or those with limited social media presence
Adapting offers of goods and services as well as prices to align with individual consumer demand	Can be used to deduce sensitive personal data from non-sensitive personal data, with a reasonable degree of certainty
Improvements in medicine, education, healthcare and transportation	Unjustifiable deprivation of services or goods
Provide access to credit using different methods to traditional credit-scoring	Risk of data broking industry being set up to use information for their own commercial interest without individual's knowledge
Can provide more consistency in the decision-making process	Using profiling techniques can jeopardise data accuracy

14

You and I know that machine learning, profiling, taking historic data sets, and using that as a basis to make an inference to predict something in the future, that uses a level of complexity that those simplified audit trails just don't begin to grapple with.

[moves to next slide]

Computers
can say 'no'



Government decision maker: who 'decides'?

Who is the decision-maker?

To whom has authority to decide been delegated?

- Authorised decision-maker?
- Policy maker?
- Computer programmer?
- Computer itself (AI)?

s 495A(1) of the *Migration Act 1958* (Cth):

The Minister may arrange for the use, under the Minister's control, of computer programs for any purposes for which the Minister may, or must, under the designated migration law:

- (a) make a decision; or
- (b) exercise any power, or comply with any obligation; or
- (c) do anything else related to making a decision, exercising a power, or complying with an obligation.

Similarly see s 6A of the *Social Security (Administration) Act 1999* (Cth).

15

And it raises the issue: who is the decision-maker? It's not a nameless bureaucrat - who actually can be named and can be found. What if it's delegated to the computer itself?

Now legislation in some respects already lets this happen, but without even anticipating the complexities in it.

[moves to next slide]

Is it ok for
computers
to say no?



Is concept of delegation appropriate in this context?

- Can a computer program act independently of its programmer or relevant government agency, just like a human delegate?
- What if computer determines some, but not all, elements of administrative decision?
- Should determination of those elements be treated as subject of separate decisions from those elements determined by human decision-maker?
- What happens when technology allows a shift from use of decision *support* tools, where intelligent systems *replace* human discretion and judgment?

16

So is delegated decision-making even appropriate? Should, as a question of the rule of law, should we be permitting decisions to be delegated to an artificially intelligent being?

[moves to next slide]

Managing the 'Black Box' Problem



Testing the *Better Practice Guide* Audit Trail requirement against the 'black box' problem:

- Is the predictive algorithm a (trade) secret?
- Was implementation secret?
- Is process (precisely) described?
- Is the process of reasoning by an autonomous system unknowable?
- Has the human middleman been rendered obsolete?

17

And the problem becomes the Black Box problem. I used jingoistic terms in the paper originally because trying to get judges to think, "*Oh, this is new. I don't know about that. I don't like not knowing about things. I better read up on this.*" Found a great *Royal Society* publication, told them about that.

But one of the issues might be for example - if you're using a predictive algorithm which is a trade secret - maybe it shouldn't be used by government. Maybe it should be only open source software that should be used.

[moves to next slide]

Sustaining public law values



- Should algorithms be accountable? Can opacity be made transparent? Do we need new rights to sustain public law values (transparency, accountability, fairness)?
- Right to explanation?
- Right to be forgotten?
- Right not to be the subject of automated decision-making?
- Right of technological due process?
- Do we need new regulators/supervisors?
 - Federal Robotics Commission (Chopra)
 - AI Policy Council (AHRC)

18

And so then just taking a step back. Public law values are just an aspect to the rule of law. Some of those values are transparency, accountability and fairness.

Transparency- how is the decision reached? What were the integers? How did they interact? What was given, what was considered relevant? What was a not relevant consideration?

Accountable - Who made the decision? When, why, how?

And fairness, considering who is the subject of the decision.

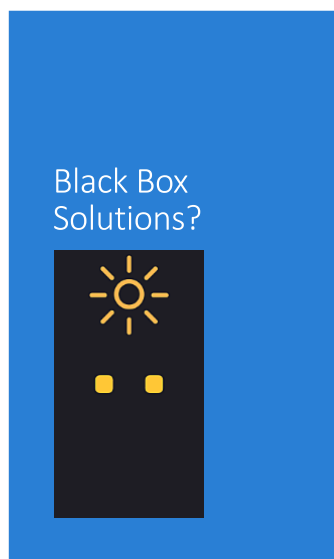
And so there's a whole debate raging primarily at the moment in Australia at the academic level, which is why my paper was trying to shove it out into more actors in the legal system.

But in the UK and in Europe obviously much more advanced because of the GDPR.

So should there be additional layers of rights which makes all of this change? Should we have a right to have an explanation? Should we say, I should have a right not to be subjected to an automated decision. Should there be a right of technological due process? That comes from the US literature, which is very much from the Bill of Rights - US process oriented jurisprudence.

Or do we need new regulators and supervisors? So in the US there is this idea of a Federal Robotics Commission. And here in Australia, the Australian Human Rights Commission and the World Economic Forum have released a white paper, *Artificial Intelligence Governance and Leadership*, which asks, should we have a responsible innovation organisation, and what should it look like? Do we need to have an AI Policy Council? And that's really important because it's trying to draw all these disparate ideas together.

[moves through next slide]



- Establishment of ethical framework
 - UK House of Lords Select Committee on Artificial Intelligence "AI in the UK: ready Willing and Able?" 2018 (urged adoption of cross sector code of 5 principles)
 - European Commission High-Level Expert Group (HLEG) on Artificial Intelligence:
 - Beneficial to humanity
 - Not infringe on privacy or undermine security
 - Protect and enhance our autonomy & ability to take decisions and choose between alternatives
 - Promote prosperity & solidarity in fight against inequality, discrimination and unfairness –innovation inclusive, promote diversity & tolerance
 - Transparent (how do AI systems work) and accountable (why do AI systems reach conclusions they do)
- Should ethics be "baked in" the code by designers?

19

I was saying to my son who's studying science here at the University, *why was I coming at all today?* I said, "*We need to have more cross-pollination of thinking.*"

And just as you've got silos in each of the experts that have come today, lawyers, and judges, and policy-makers at that level suffer from great technical illiteracy.

And the rapidity with which technological innovation is happening just cannot take with it people like lawyers and judges, and even politicians who are making laws, unless you embrace them and take them with you now. Because we'll just have intractable difficulties ultimately.

[moves to final slide]

Concluding remarks: a call to action

Ethical issues almost always directly translate into concrete legal challenges — or they give rise to difficult collateral legal problems.

A challenge for policy-makers is to develop an approach to assessing the ethical risks associated with technology assisted decision-making systems that will be simple & pragmatic, easily incorporated as part of a standard risk assessment exercise, which can be undertaken during a system's design phase.

Intelligent robotics is shaping up to be the next transformative technology of our times. It merits systematic reassessment and changes to law, institutions, and the legal profession and academy.

Critical evaluation is necessary to avoid addressing policy questions in a piecemeal fashion, with increasingly poor outcomes, and slow accrual of knowledge.

The legal profession, judiciary and academy should ensure they are well placed to support an informed public debate about what we want algorithms, machine learning, and AI to do, and how the benefits can best be distributed.

20

So, my final call. This conference is about ethics. But may I remind you that ethical issues almost always directly translate into concrete legal challenges, or they give rise at least to difficult collateral legal problems. Let's make sure all of the academy is engaged in helping find solutions. Thank you.

[Question & Answer]

[Question from the floor concerning re recent advances in machine learning, eg counter factual explanations, can quiz machine learning systems. We are learning more about how humans work, but we actually know much more about the computers than we do about the humans.]

So I don't disagree with you. The administrative decision-maker, which is the context in which I was approaching this, is *assumed* to be a rational actor, and it's when they're not being a rational actor, by for example taking into account an irrelevant consideration, or failing to take into account a relevant consideration.

But you're absolutely right. And in a sense it's an acknowledgement that discretionary decisions are so hard to really get behind when they're done by a human, that there's in effect an assumption that we can't even imagine a sufficiently intelligent artificial being, being able to be at least at that level where all these considerations can be taken into account.

Now that's a clash both of a philosophical nature and of a technical nature. But I suppose your proposition is to say, well, machine learning has advanced substantially maybe to a point where it could substitute for a human decision-maker.

Or, another way to look at it is that each is as bad as each other. But we have to have decisions made. There has to be implementation of government policy. There has to be ways in which we enhance them. And ultimately it may be that for the rule of law, we have to make choices. We have to choose between which values will we elevate. Will we elevate transparency or not? Or we're going to say, we only require transparency to a certain level. We're prepared to take the rest on trust.

[Thanks again.]

Thank you.