

Avoiding (and investigating) automated maladministration

Paul Miller PSM, NSW Ombudsman

(This is an edited version of a presentation given at the 13th National Investigations Symposium, 25 May 2023, Four Seasons Hotel, Sydney)

Almost 20 years ago now, in 2006, the NSW Ombudsman investigated the use by NSW Police Force (NSWPF) of a *non-human intelligence tool* that had been procured and trained to detect certain criminal offences.

The tool was estimated to have a technical capability in the field of criminal offence detection that was at least 40 times more powerful than humans.

Over a two-year period, the tool was used in over 470 operations.

It profiled hundreds of thousands of people, and from them detected over 10,000 potential criminals.

Most of the people identified by the tool were then subject to Police action.

I hope by now that you have figured out what ‘non-human intelligence tools’ I am talking about...

Yes, back in 2006 the NSW Ombudsman looked into the use of drug detection dogs by the NSWPF.

Dogs are estimated to have an ability to smell – let’s call that olfactory intelligence – that is around 40 times superior to humans.

In our report, we noted that during the review period 17 dogs made 10,211 indications.

That number is of course a fraction of the total number of people ‘screened’ by the dogs.¹

Of those detected, most were then subject to a search by Police.

This accorded with NSWPF policy, which stated that an indication by a drug detection dog, of itself, gives police reasonable suspicion to search a person.

Why, you might ask, am I talking about dogs when I am supposed to be talking about automated decision-making (**ADM**) and artificial intelligence (**AI**)?

Well, I'm going to leave you to ponder that as I proceed with my presentation.

But let give me you a bit of hint by pointing out just some of the issues that were raised by the Ombudsman when investigating the use of drug detection dogs:

(a) Questions of lawful authority

Police in NSW had started using drug detection dogs many years prior to the 2006 Ombudsman report.

In 2001, a Magistrate dismissed two drug charges against a man found in possession of prohibited drugs during a drug dog operation. The Magistrate held in that matter that the actions of the NSWPF using dogs constituted an illegal search.²

In essence – the statute that gave Police the power to stop and search people, did not implicitly carry an authorisation to search using a dog.

As a result of that court decision, legislation was introduced to expressly authorise NSWPF to use drug detection dogs.

Now I'm going to get straight to the point here by noting that, as we'll soon see, in the field of ADM, quite often the very first question that arises is whether it's legally permissible to use an ADM tool, or whether an amendment to the statute must or should be made to authorise that.

(b) Questions of accuracy

This was a key focus of the Ombudsman's investigation, and led to its perhaps most often quoted finding, which you still see quoted in the media today.

The stated objective of drug detection dogs is to indicate people currently in possession of illicit drugs.

So how good are they at doing that?

Well, perhaps not very. The Ombudsman's report found that they were successful in doing that just 26% of the time.

That means nearly three quarters of the people searched after being indicated by a dog were found not to be in possession of any drugs.³

Of course, if a human police officer – no matter how well trained – was tasked with sniffing people in a crowd to see if they could detect those who might be in possession of drugs, I’m pretty sure the failure rate would be a lot greater.

The question is this: is the mere fact that dogs are so much better than humans at this task enough? Or should we require some higher standard of accuracy before the outputs of a tool like this can be used as a basis for branding a person as a suspected criminal and subjecting them to an invasive search?

Related to this issue of accuracy, there is quite a bit of discussion in the Ombudsman’s report about the testing and accreditation of drug detection dogs – some dogs are obviously going to be better drug-detection tools than others – and again I invite you to consider possible parallels with AI tools.

There is, for example, an interesting discussion in the report about the differences between the testing and training environments, and real-world applications – such as how dogs tested in a controlled environment translate to large crowd situations.

(c) Questions of potential bias and discriminatory impact

The Ombudsman observed that drug detection dogs indicated men far much more frequently than women.

They also more frequently indicated young people – almost half of the people indicated were under 25.

The report was unable, because the data wasn’t collected, to say much about Aboriginal status or ethnicity of the people searched.

(d) Unintended harms

Some of the observations in the Ombudsman’s report concerned the extent to which the use of drug dogs might cause unintended harms. The things looked at included whether they might have resulted in:

- the consumption of larger amounts of drugs at once instead of taking smaller amounts over a period of time,
- consuming drugs at home and then driving to venues,
- purchasing drugs from unknown sources at venues to avoid carrying drugs,
- switching to potentially more harmful drugs in the belief that they are less likely to be detected by dogs.

So, let's turn now, more explicitly, to the topic of 'avoiding and investigating automated maladministration'.

Let me very quickly outline what I mean by the two key concepts inherent in my presentation title:

- (1) automation, or ADM, and
- (2) maladministration.

First ADM.

In the interests of time, I will not spend too much time on this terminology. I'll also probably jump around using terms like ADM, AI, or algorithmic decision-making, as if they all mean the same thing, which they don't, necessarily.

But generally speaking, if you just accept that what I am talking about is the use of technologies to make or contribute to making government decisions – decisions that had previously been the exclusive domain of human beings – then we'll be pretty much on the same page as to what I mean.

I will, however, highlight two very important points, which can sometimes be overlooked:

I referred to technologies that make or contribute to making, decisions. That's important. ADM does not just mean fully automated decision-making.

Take for example a technology that performs a filtering activity.

Mobile phone camera detection technology used in NSW is a great example. Cameras placed at intersections take and then analyse photos of drivers, identifying through the use of machine learning software those images that are most likely to show a driver holding in their hand a mobile phone device – which is of course a criminal offence. The technology therefore filters out all of those images that appear not to show such an offence.

But a fine is not automatically sent when an image is selected. Rather – at least as we understand the system to work currently – every image that gets through the AI filter is then reviewed by a human, who makes the call as to whether a fine should be issued.⁴ So, it's not fully automated decision making – but it is still ADM in the sense of contributing to a decision-making process and certainly it's a kind of ADM that I am talking about.

The second point is that although some of the technologies used in ADM are referred to as Artificial “Intelligence” – the technology does not have to be very intelligent at all, and quite frequently it is not.

ADM includes some very basic rules-based programs.

A clear example of this was some of the ADM used during COVID. Border crossing rules in various states required approval before entering a state during certain times. If you wanted approval, you filled in an online form, entering your relevant details and ticking boxes, for example to indicate your reason for wanting to enter the state. In some cases, the approval would pretty much come back instantaneously – not because there was a little human who had read your application and assessed your eligibility, but rather because a very simple rule-based algorithm had determined that you’d ticked all the right boxes (so to speak).

Here's another more recent example – a simple tool was implemented to automate the billing of ambulance and related medical services. It meant that that, whenever details were entered to show that a certain ambulance or other medical service had been provided, a bill would be automatically generated and sent to the person who received the service.

Not surprisingly, those rescued from rooftops in the recent year’s flooding found it a bit insensitive when they were automatically issued a bill for their helicopter rescues.⁵

Turning now to ‘maladministration’.

Maladministration is the central concern of all Parliamentary Ombudsman, and being the Ombudsman, it is unapologetically what I’m focused on today.

Literally maladministration means wrong or bad administration, and it the opposite of conduct that accords with administrative law and principles of good administrative practice.

It is critical to note that maladministration certainly includes, but is wider, than conduct that is unlawful, and certainly it is considerably wider than the kinds of conduct that would, if challenged in a Court of Law, be overturned or declared to be beyond power.

That said, administrative conduct that is contrary to law is the first kind of maladministration.

Illegality is the start of maladministration, but it is by no means the end.

As an Ombudsman, we have powers to make findings not just when government agencies and officials have acted unlawfully, but also when they have acted unreasonably or unjustly, or in an improperly discriminatory manner, or in some way that is otherwise wrong.

So bringing these two things together, ADM and maladministration – what can we say?

Well, the first thing we can and should say – even if it is obvious – is that the use of ADM technology is not always and inherently a form of maladministration.

Indeed, in some instances the technology has the potential to improve administrative conduct and help to address some of the things that an Ombudsman might otherwise be quite concerned about – such as by improving the consistency of decision-making and mitigating the risk of individual decision-maker idiosyncrasy and bias.

However, if ADM technology is designed and used in a way that it not lawful, or does not otherwise accord with associated principles of good administrative practice, then its use could constitute or involve maladministration.

And in this regard, the main issue I want us to think about today is whether, and if so how, automated maladministration is different from any other kind of maladministration, both in terms of what maladministration might look like and – particularly for this audience – how we might go about investigating it.

So that requires us to think about what changes when ADM comes into play, and just as importantly, what doesn't change.

First, let's look at what *doesn't* change.

In November 2021, we tabled a special report in Parliament titled, 'The new machinery of government: Using machine technology in administrative decision-making'.⁶

One of the key messages we wanted to get across in that report was to emphasise what doesn't change when technologies are used in administrative decision-making.

In particular, whenever new technology is introduced, it is introduced into an existing environment, including an existing environment of legal rules and norms of good practice.

The legal environment into which public sector ADM is being introduced is the one that is governed by public administrative law – the law which controls government decision-making.

Now existing legal environments may be more or less hospitable to new technologies. There can also be uncertainty, at least initially, about how the environment will accommodate and respond. There may, at least initially, be gaps and inconsistencies.

Administrative law, as it has developed over many centuries but rapidly during the last 50 or so years, is essentially principles-based. That means that it is, generally speaking, technology agnostic.

So, while the technology used in government decision making may change, the underlying concerns and the underlying norms that underpin administrative law remain unchanged.

This doesn't mean that the laws won't change, but the core principles and values about which the law is concerned are going to be familiar.

Let's consider those principles briefly.

For simplicity, we group the requirements for good decision-making in our 2021 report as follows:

- proper authorisation,
- appropriate procedures,
- appropriate assessment,
- adequate documentation.

They are pretty self-explanatory, even common sensical, and each of these are addressed in some detail in our report, and I won't go through them in any detail now, but I will refer back to a couple of them to highlight how they might apply in the context of ADM.

The first principle is a pretty obvious one: there must be legal power to make the relevant decision.

This is where I need to very briefly mention Robodebt.⁷

What occurred in Robodebt would have been just as unlawful had it been undertaken without any automation, and the unlawful elements of the Robodebt scheme weren't dependant on the presence of automation. Rather, the core issue is that Robodebt involved doing something that there just

wasn't legal power to do under the relevant legislation – to generate a debt by way of income averaging.

Robodebt shows that familiar problems can arise under familiar laws and principles, but in the context of using new data-enabled technologies.

But it's also the case that ADM is going to give rise to new problems, albeit they will still need to be considered against the already well-established and familiar framework of laws and principles that control agency conduct.

The point here is not whether or not those principles still apply – they most definitely do – but the issue is considering how they will apply in this new particular context of ADM.

I will refer here to our investigation in relation to Revenue NSW's use of ADM in garnishee orders. I won't go into detail, but the nub of that case was as follows:

There is a well-established principle of law that, if Parliament gives a particular person the function of making a discretionary decision – then that discretion cannot be fettered or abdicated.

A public servant can't, for example, have someone else dictating to them what decision they should make.

Nor can a public service agency adopt a policy that says, without exception, in every case we are going to make this or that decision.

In other words, if you're given discretion you have to be prepared to exercise it.

Now Revenue NSW adopted ADM technology for its garnishee system, under which it is able to sweep money directly out of people's bank accounts to recover unpaid fines debts.

Relevantly, the legal decision to do so involves some discretionary aspects.

Initially Revenue's NSW machine was automatically issuing the garnishee orders. Following our concerns, they changed the process so that a relevant human being would, at the end of each day hit "go" to issue the orders.

We obtained a legal opinion from Senior Counsel that both before *and after* this change, Revenue NSW's conduct was unlawful, as it was not in compliance with the legislation.

It seems clear why it was not lawful before there was a human 'in the loop'. But even after a human had been 'put on top', Counsel still considered the system to be operating unlawfully. In essence, it was not consistent with having a discretionary function for a human to be simply adopting the output

of a machine, without engaging themselves in an active mental process of deciding whether to issue the order – so that the decision could be said to be truly their decision and not just what the machine told them to do.⁸

It's a neat example, I think, of a very familiar rule regarding discretion being applied to a new situation presented by ADM technologies.

There are other examples that raise similar questions as to how familiar law and familiar principles might apply to different technology:

- Right to be heard

Take the right to be heard. It is a long-standing principle of procedural fairness that, generally speaking, a person who will be materially and adversely affected by an administrative decision should have a right to be heard before the decision is made. And for that right to be heard to mean anything, the person has a right to be told the case against them, so they can understand what they are responding to.

Now in the context of ADM, one question is whether that familiar right means that a person has a right to be told whether and how ADM technology has been involved in the making of the decision.

In our report we suggest that it does – we say that, in order to genuinely have a right to be heard, a person needs to understand the proposed decision and how it was made, including if it is the result of some output from a machine.

- Algorithmic bias

I am not going to talk in any detail about algorithmic bias. But there are three quick points to leave with you – especially for those of you who are lawyers – in terms of considering how the existing principles of administrative law will apply where systems might exhibit algorithmic bias.

First, the traditional *rule against bias* might not be the key one to consider – because that is primarily concerned with prejudice or bias in the mind of the individual decision-makers.

Second, laws prohibiting discrimination on certain protected grounds – like sex and race will clearly be relevant.

And third – and of most interest I think – the familiar rules concerning proper decision-making, such as the requirement to only have regard to relevant and to not have regard to irrelevant considerations, will likely affect the legality of using systems to make decisions when those

systems may exhibit algorithmic bias. Likewise, the use of a system that exhibits algorithmic bias may raise questions around compliance with the requirement that decision-making be reasonable – and I mean reasonable both in the legal sense in cases like *Wednesbury* and *Li*,⁹ but also the broader sense of reasonable as used in the Ombudsman Act.¹⁰

- Obligation to provide reasons

I will say something quickly about the obligation to provide reasons.

Although it is not always a legal requirement, it is a norm of good administrative practice that if a decision is made – particularly one that adversely affects the rights or interests of an individual – reasons should be able to be, and should be given.

‘Computer says no’ is unlikely to be a reason.

The question of what reasons can and should be provided when using an ADM systems – particularly those that use machine learning technology – is a challenging one, and continues to be a matter of some debate, but again the underlying principle that people deserve a meaningful reason they can understand is familiar.

Let’s turn to some of the things that do change with ADM...

(a) Likelihood of error

One reason why errors, or at least certain types of errors, might be more likely with ADM is because of the challenge of translating language. Laws are currently drafted in natural language, while coding is more precise with a narrower vocabulary.

Furthermore, those involved in the coding of an ADM system are generally not the kinds of people who have experience and expertise in interpreting and applying legislation.

Now this may change over time, but for now it seems that what we might call translation risks remains a considerable challenge.

(b) Consequences of error

A key advantage of machines is their ability to process high volumes of data and cases at very high speed. This also means though, that any errors will be replicated at a rate exceeding that of any individual human administrator.

Consequently, the number of people adversely affected by a single error may be much more substantial.

(c) Speed

The speed and scale of technology advancement is changing.

It is worth noting that generative AI applications such as ChatGPT were not even on the radar publicly so far as we can tell when we published our report less than 18 months ago. No doubt, it was already in development and since ChatGPT was released, there has been a proliferation of generative AI use cases – a number of which are freely available and come with significant risk. This is the challenge of keeping up.

(d) Detection of error

The detection of error may be a challenge. Detecting errors in ADM may call for quality assurance capabilities that administrators currently do not possess.

Those affected by erroneous decisions – particularly if they are already vulnerable – may be less able to identify or effectively challenge a machine error.

In both Robodebt and the Revenue NSW examples, the initial complaints were not about unlawful ADM – people were just concerned by the outcome of the process. In most cases they would have had no idea that ADM was even being used.

(e) Reversibility and rectification of errors

If a human decision maker makes an error, their conduct can usually be easily corrected for future decisions.

Even a systemic error that is reflected generally in agency policy can generally be remedied quickly and effectively.

However, if there is an error in ADM technology, fixing the error can be difficult, costly and time consuming – and a tweak here or there to fix one error may run the risk of creating others.

In our 2021 report we included an example relating to Transport for NSW and its system DRIVES which is used for a range of functions including driver licence suspensions. DRIVES was programmed in such a way that a different process is followed depending on how long the driver's licence had left until expiry at the time of the suspension:

- if there are 35 or more days left to expiry, a notice will be automatically issued.
- if there are fewer than 35 days left to expiry, no notice is issued. Instead, when the driver applies to renew their licence they will be denied a licence and given a licence suspension notice.

It appears that those coding the system made certain assumptions, including that any driver whose licence was expiring would apply promptly for a new licence. We handled a complaint from a driver whose licence was due to expire, but did not immediately apply to renew their license because they knew that they would need to serve a licence suspension period due to an accrual of demerit points. This meant, however, that a notice of suspension was never issued until the person (many months later) did apply for a new licence – and consequently it was only then that the suspension period commenced . This meant that the driver was unable to drive for a much longer than necessary period of time.

While TfNSW committed to fixing the glitch, it said it would not be possible to do so until the next scheduled system update, and that it would need to consider what interim measures it could put in place.¹¹

(f) Dynamism and creeping complacency

Another challenge is dynamism, which sounds like a good thing, but just means that the system will change over time.

Obviously, every time there is some change in the law governing the relevant function, the system will need to be modified, and over many years a system may end up looking quite different and certainly less elegant than its original design.

A related challenge is the potential for change in the relationship between the ADM and the humans who are working with it. Even if initially the system is designed so that humans are actively playing their role in a decision-making system, there's a well-documented tendency toward technology complacency – over time people just tend to place increasing reliance and confidence in technological outputs.

This means that even a system that is designed to be lawful – for example with humans properly exercising discretion – might in practice degrade, and cease to operate lawfully if those or future decision-makers cease in practice to genuinely perform that task.

(g) Cost

Another challenge is the expense associated with making changes and even small tweaks to ADM systems.

I'll just give you a sense of what we are talking about here. During the election, the now Government made an election promise that drivers with demerits points will be eligible to have 1 demerit point removed if they do not incur any further penalties over a 12-month period. This was promoted as an incentive to safe driving.

It sounds simple enough.

Like other election promises, this one was costed by the NSW Parliamentary Budget Office.

The proposal was estimated to cost \$5.66 million. The estimate includes:

- \$2.81 million to implement and test changes to systems, and
- \$2.85 million in staff resources representing an increase of 21 full time staff to support the policy including by responding to customer inquiries and manually investigating individual cases.¹²

(h) Ownership and control

Finally, I want to note that many technologies are owned and controlled externally to agencies. This can have a range of impacts on an agency's:

- ability to make changes the system,
- essential understanding of how the technology works,
- ability to provide information about the operation of the system for example to members of the public.

What about the new challenges raised by ADM specifically for Ombudsman and those who will be involved in investigating maladministration? There are a few obvious ones:

(a) System complexity

There are different phases of ADM systems from design to implementation and ongoing monitoring and review. In many cases, we need to understand each of those phases in detail in order to understand how the system works, and where there may be possible maladministration.

Consider again Robodebt – which as far as ADM goes, is relatively simple.

But consider how much work the Robodebt Royal Commission has undertaken just to understand how the system was implemented and works.

Documenting roles of people and roles of machine becomes more complex with handoffs between the two.

The Royal Commission sought expert advice from Deloitte in that case, and the process flow charts they created are all available on the Commission's website.¹³

(b) Lawfulness

As I said before, a core question for maladministration investigators will be whether agency use of ADM is lawful.

A finding that an agency has acted unlawfully is a serious claim to make, and unsurprisingly agencies tend not to like it when we do.

It's not something an ombudsman would do lightly or without being pretty sure we are on firm ground.

However, particularly in this context, there is a challenge because:

- our findings are not authoritative, and it can't be known whether a court would come to the same conclusion, and
- issues relating to public sector use of ADM is still an emerging area, and there is very little Black Letter or case law to guide us.

We can't go to court for an advisory ruling, and so in the Revenue NSW matter we obtained eminent Senior Counsel advice on the legality of the system.

(c) Visibility and transparency

Another challenge is visibility. In our 2021 report to Parliament, a key observation was that NSW agencies do not have an obligation to routinely disclose when they are using ADM.

It follows that we do not know how many agencies are using, or developing, ADM to assist them in the exercise of their statutory functions. If we, and complainants, do not know when ADM is being used, that is major obstacle to effective oversight by bodies such as ours.

I would also add here transparency, as distinct from mere visibility. It is by now well-documented that many technologies are a black box in

terms of how exactly the algorithms and models work. In this case, it can be difficult if not impossible to detect issues such as bias – which leads me to the challenge of capability.

(d) Capability

Detecting bias and assessing accuracy will require in some cases, specialist skills and knowledge.

Integrity bodies face this capability challenge now when we investigate ADM systems. We can see the need for:

- internal capability to ensure staff are alive to identifying and scrutinising ADM, and
- technical capability to understand an ADM systems' technical operation where required – this could either be embedded internally or sourced externally.

Ideally a multidisciplinary team would investigate an ADM system to consider it from a range of perspectives including legal, policy, user and technical.

One of the interesting points here concerns external capability.

I referred a moment ago to our seeking a legal opinion about the Revenue NSW garnishee system. We went to a Senior Counsel, who is eminent in the field of administrative law. That person is not a judge and is not infallible and it is possible their opinion will differ from a future court if the issue were to come before it. But, that is unlikely, and the clear structure and hierarchy of the legal profession means that we can reasonably rely on the opinion we have received.

At the moment, while there are organisations which are developing capabilities in the space of analysing AI systems – for example on detecting algorithmic bias – it may be less clear who we should go to and who we can reasonably rely on for expert opinion on such matters.

In some respects, this is not a unique situation for investigators – who often need to become instant experts themselves – but I think the fact this is still so new means that related fields of forensic accounting, auditing etc are only still developing.

There is, at present, also no guidance in Australia for choosing a credible expert or body to perform something like an algorithmic audit.

(e) Emerging standards

This leads me to the next point, which is that standards of conduct in this area are only just starting to be developed.

What I mean is this – the role of an ombudsman would be more straightforward, and investigation could be more streamlined if there were in place very clear rules about what an agency must do when it implements and uses an ADM system – for example, if there were standards that said:

- it must commission an independent algorithmic audit by an expert auditor accredited for that purpose, and
- it must obtain a comprehensive external legal certification that the system is compliant with the laws governing the relevant function.

The more well-developed, comprehensive and well-accepted such standards are, then the more straightforward in many ways will be the role of a maladministration investigator.

It would mean that we could look, at least as a first step, simply at whether those requirements were met.

Rather than diving straight in to ask technically very difficult questions like: ‘Is this ADM infected by algorithmic bias?’ we could start – and possibly end – by asking a range of questions in relation to the design, implementation and operation of the system. For example, we would ask: ‘was this system properly tested for algorithmic bias, in accordance with the requirements of the relevant standards?’.

Internationally, there are a range of AI and ADM governance initiatives that anticipate requirements for the assessment and audit of a systems’ conformity with established requirements. These may provide some guidance for us. Currently the European Parliament is considering an ‘AI Act’ and we will be interested to see whether and how that legislation clarifies how an audit should be conducted and by whom.

Standards for the development and implementation of AI by Government agencies are also beginning to develop here. NSW has adopted an ‘AI Assurance Framework’. Since March 2022, it is mandatory for agencies to apply the framework to projects that contain or use non-off-the-shelf AI.¹⁴ However, the completed framework assessment must be submitted for review to the AI Review Body in the following circumstances: if the project uses AI and costs more than \$5million or was funded from the State’s Digital Restart Fund or; if the

project uses AI and mid-range or higher risks (according to the framework) remain present after mitigations.

The Framework's aims are fairly modest at this stage – analysing and documenting a project's AI risks.

Nevertheless, a failure by an agency to properly consider and apply the framework may be an important conduct issue for consideration in any maladministration investigation that subsequently arises in respect of an agency's use of an AI system.

Now let me turn explicitly to the 'avoiding maladministration' part of my presentation.

In our report, we distilled 5 key proactive steps agencies should take when introducing or reviewing their use of ADM.

- Step 1: Assemble the right team

The first is about assembling the right team, which must involve lawyers – the statute is the source of the power and agencies need people expert in that as well as policymakers, and operational and technical experts.

- Step 2: Determine the role of staff at the outset

Deciding how far a process can be automated is not an easy question. It needs to be assessed in the context of the agency's functions and legislation. Merely placing a human on top of a process may not be sufficient to properly authorise automated decision-making.

- Step 3: Ensure transparency

We recommended agencies identify early in the project how they will be transparent about their ADM use including providing reasons for decisions made using ADM where required.

- Step 4: Test early and often

We highlighted that just like other tools that support administrative decision-making, ADM systems need to be tested before going live and at regular stages once it's in operation to ensure decisions are legal, accurate and unbiased.

- Step 5: Consider legislative amendment

Finally, we recommended that agencies consider seeking legislative amendment where necessary or prudent if it might otherwise be legally risky to proceed with ADM.

Seeking express legislative authorisation for the use of ADM not only reduces the risks for agencies, it also gives Parliament and the public visibility of what is being proposed, and an opportunity to consider what other regulation of the technology may be required.

Let's get back to our furry friends where we started.

Back when the Ombudsman examined the use of drug detection dogs, the methodology used looked something like this:

- analysis of records kept by police on their use of drug dogs,
- directly observing police using drug dogs,
- reviewing court documents about the warrants for drug dog operations,
- reviewing transcripts and judgments of cases where charges were brought following drug dog detection,
- consulting with a range of community groups and police officers of various ranks, and
- examining complaints about police utilising drug dogs.

It was also useful, no doubt, to have a fairly good understanding about dogs generally, as well as obtaining some information about the proper training and handling of dogs.

What I think is interesting is what's not on this list – at no point did the Ombudsman engage expert scientists or veterinarians to try to understand the underlying mechanism of a dog's nose to work out exactly how exactly it might be detecting drugs.

Now investigating AI may be different, depending on the circumstances, and it may be necessary for investigators to really 'get under the hood'.

And we shouldn't take the animal intelligence/artificial intelligence analogy too far.

However, there is one key aspect of both cases that I think is the same and if there is one point to take away from what I have said today it is this: when an ombudsman is investigating maladministration – we are always, and without exception, really investigating the conduct of people.

When we looked at drug detection dogs, we were concerned about the conduct of Police – their decisions and actions: how they used drug dogs, how they trained them, how they relied on them, and so on.

We weren't investigating the dogs. If a dog indicated someone who wasn't in fact in possession of any drugs, we don't say that the dog has engaged in maladministration. The question is whether Police were wrong to rely on the dog's indication.

It's the same thing with AI. We are not investigating the AI as such. We are investigating the people who made the decision to design it, to test it, to deploy it, to train it, to use it, to consider its outputs, to rely on those outputs.

And when we look at those people's conduct – the touchstones we turn to are the familiar and long-standing principles of administrative law and norms of good administrative practice – asking ourselves whether the conduct is lawful, reasonable, non-discriminatory, and just.

[END]

¹ The report noted that no records of the actual number of people screened were available, 'Review of the *Police Powers (Drug Detection Dogs) Act 2001*' *NSW Ombudsman* (Report) <[Review of the Police Powers Drug Detection Dogs Act Chapter 1-7.indd \(nsw.gov.au\)](#)>; <[Review of the Police Powers Drug Detection Dogs Act Chapter 1-7.indd \(nsw.gov.au\)](#)>

² See *Police v Darby* referred to in report at n 1, section 2.4.2.

³ The NSWPF have consistently taken the position that this does not mean that drug dogs have a 75% failure rate – for example perhaps sometimes the dogs are indicating *prior* use of the drug (and sometimes it was the case that people admitted to prior use, even if no drugs were actually found on them).

⁴ See 'The new machinery of government: Using machine technology in administrative decision-making' *NSW Ombudsman* (Report) <[The new machinery of government: using machine technology in administrative decision-making \(nsw.gov.au\)](#)> p 16

⁵ It was reported that the billing process was automated and that the invoices would be waived in this case. See 'NSW Health apologises for billing Eugowra flood victims for helicopter rescues' *The Guardian* (Article) <[NSW Health apologises for billing Eugowra flood victims for helicopter rescues | New South Wales | The Guardian](#)>

⁶ See n 4

⁷ See Robodebt case study in report n 4, p 26

⁸ See 'The new machinery of government: Using machine technology in administrative decision-making – Annexure A' *NSW Ombudsman* (Report) <[The new machinery of government: using machine technology in administrative decision-making - Annexure A \(nsw.gov.au\)](#)>

⁹ *Associated Provincial Picture Houses Ltd v Wednesbury Corporation* [1948] 1 KB 223; *Minister for Immigration and Citizenship v Li* [2013] HCA 18

¹⁰ s 26 *Ombudsman Act 1974* (NSW)

¹¹ See n 4, p 43

¹² 'Demerit points incentive' *NSW Parliamentary Budget Office* (Webpage) <[C1326 - Costing - Demerit points incentive.PDF \(nsw.gov.au\)](#)>

¹³ 'Deloitte – Robodebt Process Maps' *Royal Commission into the Robodebt Scheme* (Webpage) <[Deloitte – Robodebt Process Maps | Royal Commission into the Robodebt Scheme](#)>

¹⁴ 'NSW Artificial Intelligence Assurance Framework' *Digital.NSW* (Webpage) <[NSW Artificial Intelligence Assurance Framework | Digital.NSW](#)>