Public Law Project

Experiments in Automating Immigration Systems: A discussion

Tatiana Kazim

MARCH 2022

The Public Law Project (PLP) is an independent national legal charity. Our mission is to improve public decision making and facilitate access to justice. We work through a combination of research and policy work, training and conferences, and providing second-tier support and legal casework including public interest litigation.

Our strategic objectives are to:

- Uphold the Rule of Law
- Ensure fair systems
- Improve access to justice

www.publiclawproject.org.uk

Tatiana Kazim is PLP's Research Fellow in Public Law and Technology. She holds a Bachelor of Civil Law and an undergraduate law degree from the University of Oxford, and a politics and philosophy degree from the LSE.

Introduction

.

On 27 January 2022, Public Law Project and the University of York held a launch event for a new book by public law and technology experts Jack Maxwell and Dr Joe Tomlinson, *Experiments in Automating Immigration Systems (Experiments)*. The guest panel were:

- Book authors, **Jack Maxwell** and **Dr Joe Tomlinson**.
- **Tatiana Kazim**, a research fellow at Public Law Project, who highlighted problems that can arise even in relation to simple algorithms, including the problem of automation bias.
- **Dr Brendan McGurk**, a barrister at Monckton Chambers, who compared the group-level risk profiling undertaken by the visa streaming tool with individualised risk profiling in the insurance sector.
- **Divij Joshi**, a doctoral researcher at University College London, who argued for a fundamental change in the way we understand and respond to algorithmic technologies: one framed in the language of human rights.

Each speaker brought a unique professional perspective and this made for a rich discussion of the issues arising from *Experiments*. Their arguments are summarised in this report.

Why write this book?

Governments around the world are increasingly embracing automated decision making. The potential benefits of this are well-known: faster, cheaper, more accurate, more consistent decision making.

But we also know that, when government automation goes wrong, it can have terrible consequences. These consequences are often visited on marginalised groups: prisoners, welfare recipients and – as emphasised in *Experiments* – migrants. It's essential that we understand the problems with government automation and how to address them. As **Jack Maxwell** explained, this provided the motivation for writing the book.

Experiments argues that a precautionary approach to automated decision making is essential to ensure we all benefit from government automation without exposing people – especially marginalised groups – to unacceptable risks.

It does this by exploring the impact of flawed systems through three case studies:

- A voice recognition system used to detect fraud in English-language testing.
- An automated data matching system, used to process the large volume of applications to the EU Settlement Scheme.
- An algorithm used to categorise visa applications according to risk, where risk was assessed by reference to nationality, amongst other factors

When government automation goes wrong, it can have terrible consequences. These consequences are often visited on the most marginalised groups.

Experiments approaches the issues through the lens of administrative justice. Three key points can be drawn from the case studies:

- This area is at risk of becoming a 'law-free zone'. We need to think about how existing rules can be used to regulate automated decision making, and about what new rules might be needed.
- Automated decision making is likely to make certain kinds of problem in administration more common. Chief among these are discrimination and a lack of communication between those making the decisions and those subject to them. The experiences of those processed in this way are likely to be distinctive. When computer makes an adverse decision about you, without any human input, this is likely to give rise to a unique sense of indignity or injustice.
- The distinctive nature of the problems with automated decision making calls

for distinctive forms of redress. Too often, redress is an afterthought. Instead, proper avenues for challenging a decision and obtaining a remedy should be considered from the outset.

Speaker insights

Automation bias: even simple automation can be seriously flawed

When people think about government automation, their minds often jump straight to the really complicated algorithms: black boxes that use machine learning.

But, as **Tatiana Kazim** noted, not all algorithms are like this. One great thing about the book is that it covers a range of different algorithms. It acknowledges that some are relatively simple, involving a human being programming in some basic rules. And it acknowledges that problems can arise even in relation to these simpler algorithms. One such problem is automation bias.

Automation bias can colour the way a human decision maker interacts with an algorithm. In 1996, psychologists Mosier and Skitka defined it as the human "tendency to use automated cues as a heuristic replacement for vigilant information seeking and processing…" This definition is still in currency over 25 years later. People have been aware of this problem for quite some time. Unpacking the definition, automation bias tends to manifest in one of two ways. First, omission errors occur where a decision maker fails to take action on a relevant piece of information, because the algorithm has not alerted them to it. Second, commission errors occur when the decision maker acts inappropriately by over-relying on the information provided by the algorithm.

It's worth noting that - because automation bias concerns the human-algorithm relationship - it can only arise in instances of partial, rather than full automation. Partial automation can take the form of a streaming or triage system, which determines the type or quality of human judgment required in a particular case. One of the key examples that Tomlinson and Maxwell discuss in their book is the 'visa streaming tool': the now defunct Home Office system which automatically categorised visa applications according to risk: low ('green'), medium ('amber'), or high ('red'). A red rating meant that the application was directed to a more senior official and subjected to a higher level of scrutiny.

Another type of partial automation is a decision support system. These systems provide information to aid the human decision maker. One example is a vulnerability matrix, which looks at arrivals to the UK and gives them a vulnerability score to inform safeguarding decisions by immigration officials. If the human decision maker is hampered by automation bias, their use of the vulnerability score could go awry in various ways. To give a hypothetical example: let's say that a human decision maker is aware that a person has a history of substance misuse. They would – in the absence of any information from the matrix – treat the person as vulnerable. But the matrix does not include substance misuse as a factor, and gives that person a low vulnerability score. The decision maker relies on the score at the expense of their own judgment and does not treat the person as vulnerable when they otherwise would have done so.

Automation bias could render a decision unlawful under a whole range of public law doctrines. As Tomlinson and Maxwell point out, it could mean there is an impermissible abdication of the official's discretion, despite the established public law doctrine that an official must not apply guidance or a policy in an over-rigid manner. Similarly, dogged reliance on an algorithm's outputs could mean that an official is effectively 'shutting their ears' to an application and unlawfully fettering their discretion.

Dogged reliance on an algorithm's outputs could mean that an official is effectively 'shutting their ears' to an application and unlawfully fettering their discretion.

A further example is the way in which inappropriate reliance on an algorithm could mean that the decision maker is taking account of irrelevant considerations. For example, the Home Office uses a 'sham detection' algorithm to process couples who have given notice to be married, where one party is a foreign national and could potentially be getting married for the purpose of gaining an immigration advantage. The algorithm is designed to determine whether a proposed marriage should be investigated; it should not determine the outcome of that investigation. However, if a decision maker knows the couple they are investigating has been 'failed' by the triage system, they may take this factor into account when determining the outcome of the investigation.

What can be done to solve this problem? There is no straightforward answer. But the

literature identifies two broad categories of mitigation: the first category focusses on the human decision maker and the second focusses on the design of the automated system.

In the first category, mitigations include training the decision maker about the level of reliance to place on the system and enhancing their sense of personal accountability. Some research suggests that the more accountable the decision maker is or feels, the more closely they will scrutinise algorithm's outputs.

In the second category, mitigations can be as simple as changing the way information displays on the screen – for example, making it less prominent. Tomlinson and Maxwell note that some public bodies avoid using 'red' or 'green' ratings, as they may exacerbate automation bias. Neutrally presented information will be better than a command.

One mitigation that sits across both categories is linked to transparency. It's true that automation bias can arise in relation to even simple algorithms as well as more complex 'black box' algorithms. But there is evidence that if an algorithm is designed so as to be transparent and explainable, and the decision maker is trained to have a good understanding of how the algorithm works – including both its strengths and limitations – the risk of automation bias is reduced. Of course, there are many other arguments in favour of transparency and explainability, but perhaps guarding against automation bias is another one to add to the list.

There are many arguments in favour of transparency and explainability, but perhaps guarding against automation bias is another one to add to the list.

Individualised risk profiles: lessons from the insurance sector

Dr Brendan McGurk observed that insurers have, for some time, been using personal data, and sometimes sensitive data, to develop individualised risk profiles. It was against the backdrop of his work in this area that he approached discussion in *Experiments* of the Home Office's 'visa streaming' tool.

A striking feature of the tool was that the criteria or benchmarks against which it

judged visa applications were related to group-level characteristics. The tool operated primarily by using nationality and nationality-related criteria. These criteria included the rate of refusal of entry, the rate of immigration breaches such as over-staying, and the rate of criminal offences committed in the UK amongst those of a particular nationality. This meant that every person of a given nationality was 'tarred with the same brush' - even if, on an individual level, that person posed no risk at all.

The algorithm used these criteria to categorise applicants as 'red', 'amber', or 'green'. A red-or amber-rated applicant would be subjected to a higher level of individualised cross-checks. For example, an amber-rating would result in checking of employment details and banking information. A red-rating could result in counter-terrorism checks or DNA checks.

The essential design problem with the tool was that applicants who were red-rated on the basis of their group-level characteristics were much more likely to be refused. The statistics on this were shocking. The tool is an illustration of the "tyranny and discrimination" of unsophisticated, uniform group-based profiling.

The visa streaming tool is an illustration of the "tyranny and discrimination" of unsophisticated, uniform group-based profiling.

This can, however, be contrasted with the more sophisticated risk-profiling that takes place in the insurance industry. Government has a legitimate interest in limiting entry to those who are "conducive to the public good". And they could, in principle, make entry subject to fairly sophisticated data processing. It was surprising, then, that the Home Office did not commission a tool that made better use of big data and smart algorithms to build more accurate, individualised risk profiles.

In 2020, the Home Office committed to redesigning the tool, following proceedings issued by The Joint Council for the Welfare of Immigrants and tech-justice organisation, Foxglove. It is highly likely that the redesigned tool will follow an insurance-based model; one that operates on a much more individualised basis.

However, individualised risk-profiling is not a panacea and gives rise to its own problems. Designers may encode indirect and potentially more invidious forms of

discrimination by using facially neutral criteria that are actually proxies for characteristics like nationality and other protected characteristics. For example, a media furore arose when postcodes were relied on as a proxy for risk in the car insurance context. This resulted in higher premiums being charged in areas with higher proportions of minority ethnic groups. There may be fair and objective justifications for the use of such proxies, but in redesigning the streaming tool it will be essential to consider whether new proxies are in fact means of doing the same job as the use of group-level characteristics.

In the future, there is likely to be a much sharper focus on design questions of this kind, including the data used to train algorithms.

Against solutionism: the limits of reforming racist technologies

Divij Joshi presented a different perspective, challenging the suggestions for reform made by previous speakers. He called for a more fundamental change in the way we think about and respond to the deployment of algorithmic technologies: one framed in the language of human rights.

Divij posed two broad questions: first, how do we understand the structural and systemic challenges that these new technologies pose to democratic values and fundamental rights? Second, how do we bring these conversations into policy regulation and governance?

Divij argued that the use of data and algorithmic technologies in immigration control has its "roots" in "racist and colonial systems". Phrenology, race science, and assumptions about poor countries and how their populations behave have a fundamental relationship with the way these technologies have developed. The streaming tool, with its risk-profiling based on nationality and use of historical data on minoritised groups, is a clear example of this. But there are parallels in many other algorithmic systems. Facial recognition technologies, for example, "draw on a history of race science". To an extent, *Experiments* critically engages with these issues. This is a positive. Acknowledging these fundamental issues is preferable to uncritical "solutionism", which can mean that superficial changes are made at the expense of deeper engagement with the structural problems.

Algorithmic systems scale and formalise existing biases, and make them more difficult to uncover. But we shouldn't reach for the easy answer that such systems must be made more transparent and technically de-biased. The problem goes deeper. For

example, an ethnographic study by Virginia Eubanks exposed the neo-liberal foundations of welfare machinery in the US. She shows how the surveillance and risk-profiling of people through algorithmic systems places them in precarious situations and constitutes a structural harm.

Further, the state tends to hold more historical data on marginalised groups. So it tends to be these groups who are subject to experiments in automation. And it's these groups who are harmed when the experimentation goes wrong. Even if the data is accurate, it is problematic for marginalised groups to be disproportionately affected by data collection, surveillance, and algorithmic processing.

Human rights gives us a language to frame our response to these problems. Some algorithmic technologies are fundamentally incompatible with how we would like our democratic societies to function, with principles of equality and non-discrimination, and with autonomy and agency. Slowly, courts across the world are beginning to recognise this. For example, the SyRI welfare fraud detection system was struck down by the Dutch courts.

Some algorithmic technologies are fundamentally incompatible with how we would like our democratic societies to function, with principles of equality and non-discrimination, and with autonomy and agency.

There is work to be done on building meaningful systems of accountability that can address deep concerns about fundamental rights and democratic values - instead of relying on superficial paperwork like Data Protection Impact Assessments.

Globally, governments are becoming more receptive to innovative regulatory options. Options include: ethics reviews, guidelines on use, banning particular technologies, and banning particular uses of technologies.

Experiments advocates for the 'precautionary principle' which has been used in the US in the context of environmental policy. In that instance, impact assessments were not just a tick box exercise; they were used as a tool to allow policy officials to engage with affected populations. An approach of this kind could be an effective way of ensuring accountability within the context of algorithmic technologies.

Looking to the future

Dr Joe Tomlinson picked out a number of points arising from the panel discussion:

i) There are a huge number of uncertainties and complexities in this space, relating to both the systems themselves and the law. *Experiments* focusses on one policy sector: immigration. The issues are replicated across housing, welfare and more. For those interested in public law and administrative justice, there is a vast job to be done. This change is equivalent to the development of the welfare state and requires a legal response on a similar scale.

ii) So many of the automated systems that we are seeing in government, and particularly in immigration, are very simple, but they can still go significantly wrong. While it is important to have a broad-ranging discussion about different types of automation, including AI systems that self-learn, the majority of existing systems are simpler than that. We should not move beyond those simple systems, or assume that they work well, when - as *Experiments* shows - they can be deeply flawed.

iii) One interesting type of system which requires more research is the 'red, amber, green' or 'RAG-rating' system. Even though this kind of system is relatively simple, questions about how it shapes decision making outcomes are complex.

iv) Automation bias is a clear example of the uncertainty we see in this area. How courts are going to respond to automation bias, in terms of the evidence base and its implications in public law, is unclear. This is one of many features of automation that has yet to be tested in the courts in the UK.

v) We are likely to see different approaches within government to the design of automated systems. There will still be risks even where government departments choose to use increasingly personalised data. The way the law wraps around these developing systems will need to be considered on an ongoing basis.

vi) Another point raised by the visa streaming example is the poor quality of data held by government. Government departments want to use automated systems but, often, they do not have the quality data needed to do it well. Sometimes there are gaps in the information picture, and sometimes the data is inaccurate. This can lead systems to malfunction.

vii) The conversation is often about the data government is drawing upon to make

automated decisions. But with the digitisation of bureaucratic systems, government is going to be in a position to more easily collect data on the impact of its decisions. The conversation about data should extend to its use in monitoring and evaluation – not just its use in making more efficient decisions. Where government fails to collect data that was readily available, public law doctrines like the *Tameside* duty of inquiry could be invoked.

viii) It is essential that the legal community's response to failures of these systems is rooted in the experiences of affected populations. Building an evidence-base showing how these systems are experienced will be a vital piece of work.

It is essential that the legal community's response to failures of these systems is rooted in the experiences of affected populations.

Conclusion

Automation has many faces. This catch-all terms covers an array of different systems. There are common themes; most notably that automated systems tend to process some of the most marginalised and structurally disadvantaged in society. But they vary in complexity, in the way they use data, in their role within the overall decision-making process and – crucially – in the types of problem to which they give rise. There can be no one-size-fits-all solution. Sometimes, there may be no solution at all; some algorithmic technologies may be beyond reform.

The scale of the challenge is daunting. *Experiments* is an excellent primer. It gives a flavour of some of the systems being used in an immigration context, each with their own costs and benefits. But it is only the beginning. There is much more work to be done. And when new systems are being developed and deployed at a pace, there is no time to waste.

Experiments advocates for the 'precautionary principle': "given the range of risks associated with automated decision-making in immigration systems, until there is further public evidence and clear data on the impact of such systems, they should be

incrementally developed and clear safeguards, including public redress processes and monitoring systems, should be in effect." It remains to be seen whether the government will adopt this approach and slow the proliferation of new and different systems. In the meantime, there is a "giant experiment" going on in which "people's lives are at stake" - especially the lives of marginalised people. To a large extent, this experiment is going on behind closed doors, out of the public eye. Public lawyers and those interested in administrative justice must respond to the problems this poses - and fast.

There is a "giant experiment" going on in which "people's lives are at stake" - especially the lives of marginalised people. To a large extent, this experiment is going on behind closed doors, out of the public eye.

Public lawyers and those interested in administrative justice must respond to the problems this poses - and fast.

Public
Law
Project

The Public Law Project (PLP) is an independent national legal charity. Our mission is to improve public decision making and facilitate access to justice. We work through a combination of research and policy work, training and conferences, and providing second-tier support and legal casework including public interest litigation.

www.publiclawproject.org.uk