Law Commission Consultation: 14th Programme of Reform

PLP's response

Section 1: About You

(1) What is your name?

Joe Tomlinson and Tatiana Kazim

(2) What is your address?

The Design Works, 93-99 Goswell Rd, London, EC1V 7EY

(3) What is your email address?

j.tomlinson@publiclawproject.org.uk

t.kazim@publiclawproject.org.uk

(4) What is your telephone number?

020 7239 9033

(5) Are you responding to this consultation in a personal capacity or on behalf of your organisation?

On behalf of Public Law Project (PLP).

Section 2: Your idea for reform

Please use this questionnaire to tell us where you think there is a significant problem with the law. We want to know what you think is wrong and what practical problems arise. Please give us as much information as you can, even if you cannot answer all the questions.

(6) In general terms, what is the problem that requires reform?

The UK government uses algorithms to make decisions in a range of areas, including tax, welfare, criminal justice, immigration, and social care. A range of different algorithms are deployed. The algorithm may be a rule-based system, where the rules are written by the person designing the system. Alternatively, it may be a statistical system, where instructions are generated from patterns in historical data; this may be done by a human statistician or by way of machine learning. Whatever type of algorithm is used, there is another variable: decision-making may be fully automated or partially automated.

PLP recognises that automated government decision-making may have various benefits: saving time; reducing cost; and improving the quality and consistency of decisions. However, there are also significant challenges. Automated decision-making can go wrong in various ways. Common issues include:

- Mismatch between a system's overall purpose and its outputs. In social domains, an algorithm may make use of proxies. For example, it may use past events to predict the future. This could lead to inaccuracy and/or unlawfulness.
- Errors in the system's outputs (false positives, false negatives, or a combination of both).
- Bias, including as a result of problems in the design or training data. If the training data is unrepresentative then it may systematically produce worse outcomes when applied to a particular group. If the training data is tainted by historical injustices then the algorithm may systematically reproduce those injustices.
- Feedback loops, meaning that the system's decisions simply reflect the decisions it has made in the past, rather than the actual state of the world.
- Inflexibility, arising from the fact that automated systems generally work by applying fixed rules uniformly across a broad range of cases. Unlike a human decision-maker, an algorithm cannot make an exception for a novel case – it can only act in accordance with its programming.
- Automation bias: a well-established psychological phenomenon whereby people put too much trust in computers. This may mean that officials over-rely on automated decision support systems and fail to exercise meaningful review of an algorithm's outputs.
- Opacity, whether intentional, or due to the complexity of the system. The latter is a particular problem when it comes to machine learning. A machine learning algorithm may be a 'black box', even to an expert.

Further, an algorithm can make many, many more decisions within a given timeframe than a single human decision-maker. Hence, the negative impact of flawed automated decision-making may be much greater.

(7) Can you give us an example of what happens in practice? For example, if you are a solicitor or barrister, you might describe how the problem affects your clients.

Documents obtained by PLP under the Freedom of Information Act 2000 indicate that the Home Office is using an automated triage system to determine whether a marriage should be investigated as a 'sham'.

The algorithm sorts couples into 'red' and 'green' categories. A 'red' rating indicates that an investigation is required to identify or rule out shame activity.

PLP is concerned that this algorithm may be discriminatory because some nationalities – including Bulgarian, Greek, Romanian, and Albanian people – seem more likely to be targeted for investigation than others.

The criteria used by the algorithm have not been disclosed (despite repeated requests made by PLP), so there is a further concern around transparency and procedural fairness.

Another concern is that decision-making at the investigation stage may be flawed due to automation bias. If an official conducting a sham marriage investigation is aware that the couple has been given a red light by the algorithm, they may be predisposed to conclude that the relationship is a sham.

This is just one example but it is, in many respects, representative of wider patterns. It is a system that risks negatively impacting the lives of people it is processing but it is unclear how the system works and whether it is compliant with relevant legal standards.

(8) To which area(s) of the law does the problem relate? We will be looking into the existing law that relates to the problem you have described. Please tell us about any court/tribunal cases, legislation, books or journal articles that relate to this problem. You may be able to tell us the name of the particular Act or case that relates to the problem.

The legal issues broadly fall across public law, equality/discrimination law, human rights law, and data protection law. More specifically, the issues raised may relate to:

- Administrative law principles:
    - Rationality (*R (Johnson) v Secretary of State for Work and Pensions* [2020] EWCA Civ 778);
    - Relevancy;
    - The proper exercise of power, and the prohibition on undue delegation and fettering;
    - Procedural fairness (*R (Eisai Ltd) v National Institute for Health and Clinical Excellence* [2008] EWCA Civ 438); and
    - Systemic fairness (*R (Howard League for Penal Reform) v Lord Chancellor [2017] EWCA Civ 244*; *R (BF Eritrea) v Secretary of State for the Home Department* [2019] EWCA Civ 872).

- The right to private life under article 8 of the European Convention on Human Rights (*R (Bridges) v Chief Constable of South Wales Police* [2019] EWHC 2341 (Admin); *NJCM v The Netherlands* C-09-550982-HA ZA 18–388 (English)).
- Duties of non-discrimination under article 14 of the European Convention on Human Rights and under the Equality Act 2010, and the public sector equality duty under section 149 of the Equality Act 2010.
- Data protection law, including the right not to be subject to a decision 'based solely on automated processing' under article 22 of the UK GDPR.

For further discussion of our research on the administrative justice implications of these systems, see: Joe Tomlinson and Jack Maxwell, *Experiments in Automating Immigration Systems* (Bristol University Press, 2021) (forthcoming).

(9) Can you give us information about how the problem is approached in other legal systems? You might have some information about how overseas legislatures have responded or how the court or tribunals approach the problem.

The potential for different approaches to this issue is evident in the EU Commission's proposed artificial intelligence (AI) regulation, adopted on 21 April 2021. The proposed regulation would include:

- A register of high risk AI systems;
- Certification indicating conformity to regulatory standards;
- A requirement that the design of high risk AI systems allows for effective human oversight; and
- A blanket ban on certain AI systems, including subliminal manipulative systems; systems which exploit vulnerabilities related to age, and physical or mental disability to distort behaviour; public sector 'social credit' systems; and real time remote biometric systems in public spaces.

National and sub-national legislatures are also now taking more proactive steps to provide frameworks for the fair use of automation in the public sector. To the best of our knowledge, there is no authoritative comparative study of different frameworks but there is an increasing diversity of approaches, many of which could provide useful comparators for developing the UK's approach.

(10) Within the United Kingdom, does the problem occur in any or all of England, Wales, Scotland or Northern Ireland?

Automated government decision-making, and the attendant problems, occurs in all of England, Wales, Scotland, and Northern Ireland. It is occurring at both central government and local/devolved levels of government.

(11) What do you think needs to be done to resolve the problem?

PLP does not have a fixed view on what needs to be done to solve the problem. However, we would welcome a review of the existing legal framework governing the role of automation in public decision-making, and options for reform. The current framework is a patchwork and is generating considerable legal uncertainty for all involved (including government), and it also gives rise to a considerable risk of unfairness experienced by individuals going without an effective remedy. A new legal framework has the potential to mitigate the risks of automated decision-making and preserve its benefits.

(12) What is the scale of the problem? This might include information about the number of people affected this year or the number of cases which were heard in a court or tribunal over a particular period.

It is difficult to estimate the full scale of the problem. The full range of circumstances in which government uses fully or partially automated decision-making is not known, let alone the details of each automated decision-making system. However, it seems that automated decision-making is being used in a wide range of contexts, including tax, welfare, criminal justice, immigration, and social care, and the negative impacts of flawed systems are likely to be far-reaching. The Cardiff Data Justice Lab has tracked many examples, as has the organisation Algorithm Watch. In our own experience, the number of instances of automation we are finding in our work is increasing rapidly.

(13) What would be the positive impacts of reform?

Potential benefits of reform could include:

- Improved transparency when it comes to automated systems, their design, any training data, and the criteria used to make decisions;
- Greater accountability and reviewability of automated public decision-making;
- Reduced risk of direct or indirect discrimination;
- Reduced risk of unlawful decisions that violate administrative law principles; and
- Increased legal certainty.

(14) If this area of the law is reformed, can you identify what the costs or other negative impacts of reform might be?

This would depend on the details of any new legal framework. Complying with new regulations could be costly or onerous for public bodies – though this is not inevitable and the extent of any such impacts could be minimised as part of the design of the legal framework.

It is also important to note that, under the existing patchwork of laws , automated decision-making is generally much more difficult to challenge than traditional decision-making approaches, so maintaining the *status quo* will come with significant costs, when compared with the cost-saving potential of legal reform. At the same time, the uncertainty generated by the existing law means that, when cases are brought, government bodies are often not confident of the legality of their systems and often lose cases, potentially incurring serious costs resulting from the need to change an automated system in response to an adverse judgment (*e.g. Secretary of State for Work and Pensions v Johnson* [2020] EWCA Civ 778).

(15) Does the problem adversely impact equality, diversity and inclusion by affecting certain groups in society, or particular areas of the country, more than others? If so, what are those groups or areas?

One of the problems with automated government decision-making is the potential for bias within the system. By its nature, this problem involves a negative impact on some groups more than others. For example, the sham marriages algorithm appears to have a greater adverse effect on particular nationalities (see above, in response to question seven). The problem of bias may also affect other automated decision-making systems.

(16) In your view, why is the independent, non-political, Law Commission the appropriate body to undertake this work, as opposed to, for example, a Government department, Parliamentary committee, or a non-Governmental organisation?

The issues involved in this field are complex in at least four ways: (1) complexity resulting from the uncertain application of the existing patchwork of laws; (2) complexity resulting from the technical aspects of automated systems; (3) complexity resulting from the variety of automated systems used; and (4) complexity resulting from the variety of policy sectors in which automated decision-making is being deployed. The presence of such complexity means, in our view, that an independent, expert-led, consultative body such as the Law Commission is ideally placed to tackle this challenge.

(17) Have you been in touch with any part of the Government (either central or local) about this problem? What did they say?

N/a

(18) Is any other organisation such as the Government or a non-Governmental group currently considering this problem? Have they considered it recently? If so, please give us the details of their investigation on this issue, and why you think the Law Commission should also look into the problem.

N/a