University of New South Wales Law Research Series

AUTOMATING GOVERNMENT DECISIONMAKING: IMPLICATIONS FOR THE RULE OF LAW

MONIKA ZALNIERIUTE, LYRIA BENNETT MOSES AND GEORGE WILLIAMS

In S. de Souza, M. Spohr (eds), Technology,
Innovation and Access to Justice: Dialogues on
the Future of Law (Edinburgh University Press, 2021)
Chapter 6
[2021] *UNSWLRS* 40

UNSW Law
UNSW Sydney NSW 2052 Australia

E: unswlrs@unsw.edu.au
W: http://www.law.unsw.edu.au/research/faculty-publications
AustLII: http://www.austlii.edu.au/au/journals/UNSWLRS/
SSRN: http://www.ssrn.com/link/UNSW-LEG.html

M. Zalnieriute, L. Bennett Moses and G. Williams, 'Automating Government Decision-Making: Implications for the Rule of Law,' in S. de Souza, M. Spohr (eds), *Technology, Innovation and Access to Justice: Dialogues on the Future of Law,* Edinburgh University Press, UK, 2021, pp. 91 – 111

CHAPTER 6

AUTOMATING GOVERNMENT DECISION-MAKING: IMPLICATIONS FOR THE RULE OF LAW[1]

Monika Zalnieriute,[*] Lyria Bennett Moses[**]
and George Williams[***]

Abstract

This chapter assesses the benefits and challenges to the rule of law posed by automation of government decision-making. We focus narrowly on aspects of the rule of law that have the widest acceptance across political and national systems, notably that it requires governance in which the law must be predictable, stable, accessible and everyone must be equal before the law. These rule of law values are applied to four case studies: automated debt-collection in Australia, data-driven risk assessment by judges in the United States, social credit scoring in China, and automated welfare in Sweden.

---

[1] This book Chapter is a shorter version of a longer discussion in M. Zalnieriute, L. Bennett Moses and G. Williams, 'Rule of Law and Automation in Government Decision-Making,' *Modern Law Review,* 2019, Vol. 82(3), pp. 425 – 255

[*] Senior Lecturer, Australian Research Council Early Career Discovery Award Fellow (project number DE210101183), Faculty of Law & Justice, UNSW Sydney; and Lead of 'AI and Law' Research Stream, Allens Hub for Technology, Law and Innovation, UNSW Sydney. m.zalnieriute@unsw.edu.au.

[**] Director, Allens Hub for Technology, Law and Innovation, Professor, Faculty of Law & Justice, UNSW Sydney.

[***] Deputy Vice-Chancellor, Planning and Assurance, Anthony Mason Professor and Scientia Professor, UNSW Sydney. The authors thank Nayan Bhathela for research assistance.

M. Zalnieriute, L. Bennett Moses and G. Williams, 'Automating Government Decision-Making: Implications for the Rule of Law,' in S. de Souza, M. Spohr (eds), *Technology, Innovation and Access to Justice: Dialogues on the Future of Law,* Edinburgh University Press, UK, 2021, pp. 91 – 111

A. Introduction

Automation promises to improve a wide range of processes. The introduction of controlled procedures and systems in place of human labour can enhance efficiency as well as certainty and consistency. It is thus unsurprising that automation is being embraced by the private sector in fields including pharmaceuticals, retail, banking and transport. Automation also promises like benefits to government. It has the potential to make governments – and even whole democratic systems – more accurate, more efficient and fairer. As a result, several nations have become enthusiastic adopters of automation in fields such as welfare allocation and the criminal justice system. While not a recent development, automated systems that support or replace human decision-making in government are increasingly being used.

This chapter assesses the benefits and challenges to the rule of law posed by automation of government decision-making. In this regard, reference should be made to a few short commentaries which call for more attention to be paid to the governmental context: see, for example, Mikhaylov, Esteve, and Campion 2018; Kennedy 2017; and Perry 2017. It adopts the rule of law as a standard because it is accepted worldwide as providing normative guidance on the appropriate conduct and operation of governments. The rule of law is an ubiquitous, yet elusive concept, at the heart of which lies a widely held conviction that society should be governed by law. However, in this chapter, our goal is not to provide yet another account of the rule of law (modern accounts include Lord Bingham 2007: 69; Tamanaha 2004: 2; and Gowder 2016). Instead, we critically investigate how principles of the rule of law are affected by the increasing use of two kinds of automation: human-authored pre-programmed rules (such as expert systems) and tools that derive rules from historic data to make inferences or predictions (often using machine learning).

We focus narrowly on aspects of the rule of law that have the widest acceptance across political and national systems, notably that it requires governance in which the law must be predictable, stable, accessible and everyone must be equal before the law (International Congress of Jurists 1959: para. 1). In applying these principles, our focus is upon the formal and procedural aspects of the rule of law, rather than its capacity to encompass a broader set of human rights, including free speech and privacy. Hence, we limit our analysis to the following core components: transparency and accountability; predictability and consistency; and equality before the law.

These rule of law values are applied to four case studies: automated debt-collection in Australia, data-driven risk assessment by judges in the United States, social credit scoring in China, and automated welfare in Sweden. The case studies have been selected to provide a diverse range of viewpoints from which to assess the benefits and risks to the rule of law posed by the use of automated decision-making by governments around the world. We do not provide a detailed consideration of jurisdiction-specific constitutional, administrative and statutory requirements constraining decision-making in these nations. For example, in the United States, this would include due process protections in the Administrative Procedure Act, Pub L 79-404, 60 Stat 237, 5 USC §§ 551-559. Our aim instead is to analyse developments at the conceptual level of how they impact upon the rule of law, rather than seeking to develop a detailed prescription for the design or implementation of such systems.

B. Automation of Decision-Making

Automation in government decision-making is not a new phenomenon, nor is it linked to a single technology. Automation can vary from partial to full – that is from decision support (for example, a facial recognition tool that helps humans make decisions) to human-in-the-loop (decisions are made with some human involvement), to the disappearance of humans from the decision-making process entirely.

While it is not always easy to strictly categorise the degrees of automation, in analysing the impact of automation on the rule of law, we look at two classic types. The first type, known as expert systems, is a process that follows a series of pre-programmed rules written by humans, and has been used in a variety of government contexts such as child protection and the calculation of welfare benefits since the 1980s (for example, Schuerman, et al. 1989; Sutcliffe 1989). The Robo-debt and the Swedish welfare system are more modern examples of expert systems, sometimes described as the first wave of artificial intelligence ('AI') (see generally Tyree 1989; Susskind 1987: 114-15). The second type, supervised machine learning, deploys rules that are inferred by the system from historic data. It is deployed in the judicial sentencing in US and is widely known as a 'second wave' AI (Launchbury 2017). Machine learning describes a variety of data-driven techniques that establish processes by which a system will 'learn' patterns and correlations so that it can generate predictions or reveal insights. Unlike standard statistical methods, machine learning is generally iterative (capable of continually 'learning' from new information) and capable of identifying more complex patterns in data.

Despite automating the decision-making process to varying extents, neither of the approaches to automation considered in this chapter remove humans from the process entirely. At least at this stage of technological development, most of the automation comes after humans have designed and built the system. This means that the human aspect of these technologies can never be discounted. This is apparent in each of the following case studies, which we apply as the reference point for our analysis in this chapter**.**

I. Robo-debt in Australia

Robo-debt is a nickname given by the media to a controversial program, announced by the Australian government in 2015, to calculate and collect debts owed because of welfare overpayment. The program was introduced as part of a 2015–16 Budget measure, 'Strengthening the Integrity of Welfare Payments' and a December 2015 Mid-Year Economic Fiscal Outlook announcement. The system combined data matching (possibly employing machine learning), such data matching being authorised by the *Data Matching Program (Assistance and Tax) Act 1990* (Cth), and automated assessment through the application of human-authored formulae, and automated generation of letters to welfare recipients.

Under the system, data on annual income held by the Australian Tax Office (ATO) was automatically cross-matched with income reported to the government welfare agency Centrelink. To understand how it worked, it is important to know that income is reported to the ATO as an annual figure but to Centrelink as a fortnightly figure. The first step was to check the two annualised income figures against each other. Where the ATO annual income was greater than the Centrelink annualised income, individuals were sent a letter giving them an opportunity to confirm their annual income through an online portal. Those who accessed the online portal were given an opportunity to state their fortnightly income (with evidence), whereas those who did not access the portal were assumed to earn a fortnightly figure calculated as the annual ATO figure divided by the number of weeks in a year (Commonwealth Ombudsman 2017: 1, 4). However, the letter sent to individuals did not explain that recording variation in income over the year was important to an accurate calculation of welfare entitlements (Commonwealth Ombudsman 2017: 9). Concerns about Robo-debt range from poorly worded correspondence, inaccuracy of the formula in a percentage of cases, issuing debt notices to those not owing money (Carney 2018), shifting the burden of proof (Hanks 2017: 9-11), and leaving individuals to the mercy of debt collectors.

II. Data-driven risk assessment in US sentencing decisions

In some jurisdictions in the United States, judges use an automated decision-making process called COMPAS ('Correctional Offender Management Profiling for Alternative Sanctions') that draws on historic data to infer which convicted defendants pose the highest risk of re-offending, particularly where there is a risk of violence. Reliance on COMPAS in judicial sentencing has been endorsed by a Conference of US Chief Justices (CCJ/COSCA Criminal Justice Committee 2011) and by the Supreme Court of Wisconsin in *State of Wisconsin v. Loomis*, 881 N.W.2d 749 (Wis. 2016) ('*Loomis*') where it was found to be permissible so long as the decision was not fully delegated to machine learning software. For example, a judge will still need to consider a defendant's arguments as to why other factors might impact the risk they pose (*Loomis* at [56]). The United States Supreme Court denied certiorari on 26 June 2017

Concerns have been raised that African Americans are generally more likely than whites to be given a false positive score by COMPAS (Angwin, et al. 2016). This is not necessarily because race is used as a variable in modelling relative dangerousness of the offender population; differential impact can result where race correlates with variables that are themselves correlated with risk classification. In *Loomis*, data on gender was included in the set on which the algorithm was trained, the reason being that rates of re-offending, particularly violent re-offending, differ statistically between men and women. The Supreme Court of Wisconsin held that this kind of differential treatment did not offend the defendant's due process right not to be sentenced based on his male sex. Its reason was that because men and women have different rates of recidivism, ignoring gender would 'provide less accurate results' (*Loomis* at [77] and [86]). This highlights a fundamental question about the logic employed in drawing inferences using rules derived from historic data – if the goal is to maximise predictive accuracy, does it matter from a rule of law perspective whether individuals are classified differently based on inherent characteristics?

III. Automated student welfare in Sweden

The Swedish National Board of Student Finance (CSN) manages financial aid to students in Sweden for their living costs, which includes grants and various loans (see the website of CSN, n.d.). The CSN automated rule-based decision-making system is mandated by national legislation, and the role of professional officers is to guide customers through the e-service in accordance with an ethical code (Wihlborg, Larsson, and Hedström 2016). Numerous e-services provided by CSN are partially or fully automated. For example, an e-service that allows people to apply for a reduction in repayments is used to support decision-making processes (partial automation), while all decisions on loan repayments based on income over the last two years are fully automated. The automated decision-making system combines data from CSN with publicly available information, including tax information (which is publicly available in Sweden) (Swedish Tax Agency 2016). Whenever an individual applies for a reduction, an officer enters any relevant information into the system manually before letting the automated system take over again, meaning that the system is partially automated. While it is the system that 'makes' decisions, the officers are obliged by law to take responsibility for them and to communicate the decisions to the customers by editing the default formulation and signing it.

IV. Social Credit System in China

A fourth case study of automation is the Social Credit System (*shehui xinyong tixi* – 'SCS') developed by the central government of China and implemented by 43 'demonstration cities' and districts at a local level. (A linguistic note made by Creemers [2018] is useful in this context: 'the Mandarin term "credit" [*xinyong*] carries a wider meaning than its English-language counterpart. It not only includes notions of financial ability to service debt, but is cognate with terms for sincerity, honesty, and integrity.') According to the government planning document that outlines the system, translated by Creemers, 'its inherent requirements are establishing the idea of a sincerity culture [sic], and carrying forward sincerity and traditional virtues, it uses encouragement to keep trust and constraints against breaking trust as incentive mechanisms, and its objective is raising the honest mentality and credit levels of the entire society.' (State Council 2014 in Creemers 2018 [ed.]). In accordance with such goals, the SCS provides rewards or punishments as feedback to individuals and companies, based not just on the lawfulness of actions, but on their morality, and covers economic, social and political conduct (see Creemers 2018 [ed.] for a detailed analysis of the thinking and design process behind the SCS).

From a technological perspective, the SCS resembles a straight-forward, pre-programmed rule-based system, however each of the 43 'model cities' implement the programme differently. For example, under the Rongcheng City model (Junwei 2017), everyone is assigned a base score of 1,000 points on a credit management system, which connects four governmental departments. Subsequent points are then added or deducted on the system by (human) government officials for specific behaviour, such as, for example, late payment of fines or traffic penalties. There are, in total, 150 categories of positive conduct leading to additional points on the system, and 570 categories of negative behaviour leading to point deductions for individuals. The implications of the SCS cover a wide range of economic and social repercussions. For instance, those with low social credit rating scores may not be eligible for loans and certain jobs and could be being denied the ability to travel on planes or fast trains. In contrast, those with high scores enjoy benefits such as cheaper public transport, free gym facilities and priority for waiting times in hospitals.

The SCS is still in its early stages and the Chinese government has been forming partnerships with private companies with sophisticated data analytics capacities. For example, the central government has been cooperating with Chinese tech giant Alibaba in a Sesame Credit system, which includes, among other things, an automated assessment of potential borrowers' social network contacts in calculating credit scores (Hvistendahl 2017). This means that those with low-score friends or connections will see a negative impact on their own scores because of an automated assessment (Zhong and Mozur 2018). Sesame Credit combines information from the Alibaba database with other personal information, such as individual online browsing and transaction history, tax information and traffic infringement history, to automatically the determine the trustworthiness of individuals.

C. Benefits and Challenges to the Rule of Law

I. Transparency and accountability

Automation offers many potential benefits in enhancing the transparency and accountability of governmental decision-making. Whereas a human may come up with justifications for a decision ex post that do not accurately represent why a decision was made (Nisbett and Wilson 1977: 231), a rules-based system can explain precisely how every variable was set and why each conclusion was reached. It can report back to an affected individual that the reason they were ineligible for a benefit was that they did not meet a criterion that is a requirement of a legislative or operational rule that is pre-programmed into the logic of the system. It is important to note here that such feedback is not necessarily provided for rules-based expert systems. The designer decides what the output of the system will be and whether it will include reasons for its conclusions or decisions. In the case of robo-debt, individuals were not provided with clear information as to how the debts were calculated in general or in their individual case. The opposite is true for the Swedish system, where decisions are made based on clear, public rules and a human confirms and takes responsibility for each decision.

To understand the barriers to transparency, it is helpful to understand Burrell's three 'forms of opacity' (Burrell 2016: 1). The first form is intentional secrecy, which arises when techniques are treated as a trade or state secret, or when data used in the process contains personal information which cannot be released due to privacy or data protection laws. This form of opacity can apply to systems based on rule-based logic and systems that derive rules from data using techniques such as machine learning. In the case of the Chinese Social Credit system, only limited information is made public. For example, the details of the cooperation between the central government and the private sector in the Sesame Credit system are not clear. While it is known that the system will use machine learning and behavioural analytics in calculating credit scores (Hvistendahl 2017), individuals have no means to know what information from their social network contacts was used or its precise impact on their scores (Zhong and Mozur 2018).

A government agency may also outsource the building of or licence the use of an automated system and will then be bound by contractual terms that prevent further disclosure (see Noto La

Diega 2018: 11–16 for a discussion of intellectual property rights limiting the transparency of algorithms; in the context of outsourcing, there are additional considerations [beyond non-transparency] that may have legal implications that are beyond the scope of this paper). In the case of COMPAS, Northpointe Inc (now 'equivant' (Equivant, n.d.)), which built the tool, has not publicly disclosed its methods as it considers its algorithms trade secrets (this is noted in *Loomis* at 144). While the risk assessment questionnaire and thus the input variables have been leaked (see Angwin, et al. 2016), there is insufficient information available about methods and datasets used in training. While trade secret rights may legitimately be claimed by private corporations and enforced against contracting parties who agree to confidentiality provisions, there are important questions from the perspective of the rule of law about whether secret systems can be used in government decision-making in contexts that directly affect individuals. In at least some circumstances, rule of law considerations should favour open source software.

The second form of opacity identified by Burrell (Burell 2016: 4), again potentially relevant to both kinds of automation considered here, is technical illiteracy. Here, the barrier to greater transparency is that even if information about a system is provided (such as a technique used in training a machine learning algorithm or the formal rules used in an expert system), most people will not be able to extract useful knowledge from this. A system may accordingly be transparent to a technical expert, while remaining opaque to the majority of the governed, including those affected by particular decisions. In contexts where the consequences of a decision are severe, the lack of access to expert advice in understanding and challenging a decision effectively reduces the extent to which the decision itself can be described as transparent and accountable in practice.

Whereas the second form of opacity involved limitations of expertise, the third form of opacity recognises human limitations in truly understanding or explaining the operation of complex systems. It relates specifically to machine learning and stems from the difficulty of understanding the action of a complex learning technique working on large volumes of data, even equipped with the relevant expertise (Burrell 2016: 10). Because humans reason differently to machines, they cannot always interpret the interactions among data and algorithms, even if suitably trained. This suggests that the transparency necessary for the rule of law may decrease over time as machine learning systems become more complex.

There are some possible and partial solutions to this challenge. Some researchers are working on 'explainable AI', also known as XAI, which can explain machine learning inferences in terms that can be understood by humans – for example, there is an XAI program at the Defence Advanced Research Projects Agency in the US that aims to develop machine learning systems that 'will have the ability to explain their rationale, characterise their strengths and weaknesses, and convey an understanding of how they will behave in the future.' (Gunning, n.d.). It is also possible to disclose key information about a machine learning system, such as the datasets that were used in training the system and the technique that was used. Machine learning systems can also be made transparent as to aspects of their operation. However, some machine learning techniques cannot be rendered transparent, either generally, in particular circumstances or to particular people. The three challenges identified by Burrell, taken together, mean that it will rarely be possible for public transparency as to the full operation of a machine learning process, including understanding reasons for the decision, understanding limitations in the dataset used in training (including systemic biases in the raw or 'cleaned' data), and accessing the source code of the machine learning process.

An alternative solution lies in the fact that decision-making systems only need to be transparent and accountable as a whole, which does not necessarily imply visibility of the entire operation of automated components of that system. For example, in the Swedish student welfare example, a human remains accountable for the decision, even though the logic itself is first run through an automated system. Ultimately, the success of this strategy depends on its implementation. If the human can be called on to provide independent reasons for the decision, so that the automated system is essentially a first draft, then the decision-making system as a whole is as accountable and transparent as it would have been in the absence of decision-support software. If, however, the human can rely on the output of the system as all or part of their reason for decision, then accountability for the decision remains flawed despite assurances. This goes back to the question of the degree of automation in the decision-making process and the influence of outputs over the ultimate decision. A decision-making system as a whole can be made transparent and accountable by marginalising automated components (at the cost of efficiency and other benefits) and ensuring human accountability in the traditional way or by rendering transparent and accountable those automated components.

As is evident from above, the degree of transparency inherent in an automated system is a question of human design choices. Professionally, there has been a move to the development of standards, frameworks and guidelines to ensure that decision-making and decision-support systems are ethical (for example, the Artificial Intelligence, Ethics and Society (AIES) conference, the IEEE's (Institute of Electrical and Electronics Engineers) Global Initiative on Ethics of Autonomous and Intelligent Systems, the International Standards Organisation's JTC1/SC42 standardisation program, and the 'Artificial Intelligence Roadmap and Ethics Framework' project at Australia's Data61). This suggests another potential way forward for the rule of law, writing it into the language of technical specifications for decision-making and decision-support systems deployed by government. However, it is achieved, the need for greater transparency about automated decision-making software is one of the most frequently emphasised issues by both technical and legal experts (Carlson 2017: 303; Diakopoulos 2016). Citron and Pasquale have advocated for a 'technological due process', which would enable individuals to challenge automated decisions made about them (Citron and Pasquale 2014: 20). In particular, they argue that people should have a 'right to inspect, correct, and dispute inaccurate data and to know the sources (furnishers) of the data' (Citron and Pasquale 2014: 20). Furthermore, they argue that an algorithm that generates a score from this data needs to be publicly accessible – rather than secret – so that each process can be inspected (Citron and Pasquale 2014: 22). However, where automated components of systems cannot be made transparent, accountability needs to be assured by humans. Ensuring a human is responsible for independently justifying the decision and that humans are involved in appeal processes, as is the case is Sweden, is one way in which accountability can be preserved. In these situations, it will be important to ensure that such humans feel able to act independently of the outputs of the automated system. Finally, it may be the case that, because of the inherent opacity, certain decision-making by the governments, such as criminal sentencing, should not be partially or fully delegated to software with whose logic cannot be rendered transparent and comprehensible to defendants and their representatives. This ensures that factors that ought to be irrelevant in the sentencing process remain so.

II. Predictability and consistency

Automation can also improve the predictability and consistency of government decision-making. Unlike humans, computer systems cannot act with wanton disregard for the rules with which they are programmed. As such, the systems in our examples generally enhance the predictability and consistency of decision-making, even where they are otherwise problematic. The social credit system in China works as a tool of social control because people can predict the consequences of engaging in particular activities that the government wishes to discourage. Australia's robo-debt program and Sweden's social welfare system perform the same calculation for everyone.

However, automation also poses many challenges for the rule of law principles of predictability and consistency. A first challenge arises when the rule that is applied in an automated decision-making process does not correspond with statutory or common law requirements. The inconsistency in such case is not in the application of the rule in different cases, but between the rule as formulated and the rule as applied in every case. An example of such inconsistency is robo-debt. The formula failed to produce the legally correct result for many people. There is some dispute about the rate of error and how these should be characterised; approximately 20 per cent of people who received debt notices succeeded in providing additional information that demonstrated that no debt was owed (Senate Community Affairs References Committee, Parliament of Australia 2017: para. 2.88). The problem was not that there was an error rate, which also exists for decisions made by humans, but that the processes in place to manage the error were insufficient. There was no human checking of the decision to issue a debt notice. The online portal in place to deal with challenges to debt notices was also hard to use (Senate Community Affairs References Committee, Parliament of Australia 2017: para. 2.110), with human alternatives inadequate to meet the demand (Senate Community Affairs References Committee, Parliament of Australia 2017: paras. 3.98, 3.106, 3.107 and 3.119). The rate of errors also potentially exceeded the capacity of institutions designed to deal with appeals. This compares unfavourably with the automated Swedish system, where humans edit and take responsibility for each decision, with usual processes in place for appeal (CSN decisions can be appealed to the National Board of Appeal for Student Aid [Överklagandenämnden för studiestöd, 'OKS'], see OKS' [n.d.] website). The result in Australia is a far higher likelihood that the law is being misapplied in ways that are unpredictable and inconsistent.

When moving from pre-programmed rules to rules derived from data (for example, through supervised machine learning), the predictability and consistency of decision-making may be reduced. This is not because the computers are acting contrary to programming but because, like human children who 'learn', it is hard to predict the outcomes in advance and behaviour will change as 'learning' continues. Consider what is known about the COMPAS tool (which is limited due to the transparency issues discussed above). Those developing the tool did not necessarily know in advance what criteria would be found to correlate, alone or in combination, with particular behaviours (such as reoffending). The rules allocating scores to individuals were derived, likely through a supervised machine learning process, from a large set of data (namely data recording historic re-offending behaviour). The behaviour of the system is thus difficult, and sometimes impossible, for a human to predict in advance.

Machine learning raises another issue for predictability and consistency because it continues to 'learn' from new data fed into it over time. If it gives a low score to an individual, thereby contributing to a decision to grant parole, but the individual reoffends, that will be fed back into the algorithm in order to improve its predictive accuracy over time. In that way, a new individual who was relevantly 'like' the earlier false negative will have a different outcome, namely a higher risk score and lower chance of parole. This means that the system treats identically situated individuals differently over time which, as discussed below, is a problem not only for consistency but also for equality before the law.

M. Zalnieriute, L. Bennett Moses and G. Williams, 'Automating Government Decision-Making: Implications for the Rule of Law,' in S. de Souza, M. Spohr (eds), *Technology, Innovation and Access to Justice: Dialogues on the Future of Law,* Edinburgh University Press, UK, 2021, pp. 91 – 111

Moreover, the fact that COMPAS relies on variables that would not have been considered relevant by a human judge (such as whether their parents are divorced, Angwin (Angwin 2018) showing the question '[i]f you lived with both parents and they later separated, how old were you at the time?') creates an inconsistency between decisions made by judges under the law and decisions suggested by algorithmic inferencing. The lack of transparency about the data relied on in the machine learning process in a particular case, as well as opacity of the algorithm itself, makes it more difficult for judges to adjust their expectations of the tool to ensure appropriate use.

Automation according to human-crafted rules (derived from statute or judge-made law) can ensure that the correct decision is made every time and can overcome issues with human error and corruption. Rules derived from data raises more complex challenges, particularly in ensuring predictability and consistency with the 'law on the books.'. Supervised machine learning and other iterative systems also struggle with consistency over time. However, these are matters that can be controlled from the perspective of predictability and consistency, in the first case through design of the system as well as independent testing and evaluation, and in the latter by moderating continual learning. Hence, a system that combines both types of automation by using explicit programming to automate the application of a fixed rule (originally derived from data, for example through machine learning) can ensure consistency over time. Automation can thus prove beneficial for predictability and consistency, although the evidence suggests that may not be achieved in practice.

III. Equality before the law

Automation can enhance the principle of equality before the law by reducing arbitrariness in the application of law, removing bias and eliminating corruption. For instance, automation in China's social credit system could, through the use of cameras and face recognition technology, be deployed to ensure consequences apply to everyone who breaches certain rules (such as jaywalking or parking illegally) without exception. By contrast, without such automation, systems in place for minor infringements of this kind require a person to be 'caught', with the severity of the penalty often depending on the discretion and 'generosity' of the officials in question. Moreover, the enhanced consistency discussed above, particularly of the expert systems, such as the Swedish welfare or robo-debt, that give the same answer when presented with the same inputs, helps to ensure that similarly situated individuals are treated equally. These examples demonstrate how certain kinds of automation can remove the capacity for biased humans to discriminate against unfavoured groups. A properly designed system could do so by eliminating both conscious and unconscious bias by only applying criteria that are truly relevant to making the decision.

The benefits that automation can provide to equality before the law are however qualified by two main interrelated challenges. First, automation in government decision-making might compromise due process rights and the extent to which the laws apply to all equally; and second, it might undermine the extent to which people, irrespective of their status, have equal access to rights in the law.

Firstly, automation can compromise individual due process rights because it may undermine the ability of that person to influence or challenge a decision affecting them. For instance, in robo-debt, the right to review and rectify information was undermined because the letter sent to individuals by the government did not explain the importance of the income variation over the year for an accurate calculation of welfare entitlements (Kehl, Guo and Kessler 2017: 9). By contrast, the involvement of a case officer in the Swedish student welfare example enables explanation of the process and provides an immediate opportunity for those affected to rectify information or exercise a right of review. Moreover, the process is strengthened by a relatively straightforward appeal procedure to challenge the CSN decisions. For example, a student who had been prevented from joining the job market due to their disability had the initial CSN decision reversed after examination by the Swedish National Board of Appeal for Student Aid (OKS 2014). Decisions by the Board which are deemed to be of fundamental importance and in the public interest are available on OKS' website (n.d.).

Similarly, under Shanghai Municipality's SCS model, individuals have a right to know about the collection and use of their social credit information and can access and challenge the information contained in their credit reports (Shanghai Development and Reform Commission 2017: art 34; article 36 further states, '[w]here information subjects feel that there was error, omissions, and other such circumstances … they may submit an objection to the municipal Public Credit Information service center, credit service establishments, and so forth'). The municipal Public Credit Information services centre will determine whether to rectify the information within five working days of receiving the objection materials. These rights were tested in practice by Chinese citizen Liu Hu, who was blacklisted on the SCS and unable to purchase a plane ticket after he accidentally transferred the payment for a fine to the wrong account (Mistreanu 2018). After a court learned that Liu Hu had made an honest mistake, the information on his social credit report was rectified.

In the case of machine learning, lack of transparency, which is common for the reasons discussed above, is the primary reason why due process rights may be compromised. In *Loomis*, the Supreme Court of Wisconsin held that due process was preserved because a COMPAS score was only one among many other factors to be considered by the judge. (It is likely significant that the judge told Loomis at the sentencing hearing that the COMPAS score was one of multiple factors that his Honour weighed when ruling out probation and assigning a six-year prison term: '[i]n terms of weighing the various factors, I'm ruling out probation because of the seriousness of the crime and because your history, your history on supervision, and the risk assessment tools that have been utilised, suggest that you're extremely high risk to re-offend.': *Loomis* at 755.) However, the extent to which an individual decision is based on the outputs of COMPAS is difficult to assess – the Court simply added that while COMPAS cannot be determinative in sentencing decision, the risk scores can be considered a *relevant* factor in several circumstances, including: '(1) diverting low-risk prison-bound offenders to a non-prison alternative'; (2) assessing the public safety risk an offender poses and whether they can be safely and effectively supervised in the community rather than in prison; and (3) to inform decisions about the terms and conditions of probation and supervision (see *Loomis* at 767–72 [Bradley, J.], 772 [Rogennsack, C.J., concurring], and 774 [Abrahamson, J., concurring]). The Court also added that the right to review and rectify was satisfied because the defendant had a degree of control over relevant input data: he could review the accuracy of public records and offer other data directly through completion of the COMPAS questionnaire (*Loomis* at 765). However, there is a difference between the ability to review and rectify separate pieces of information which are fed into the software and the ability to review or challenge how the score is calculated.

Further, a defendant lacks an effective opportunity to challenge the idea that factors outside of his control (for example, the fact that his parents divorced when he was three, asked in the COMPAS questionnaire) influence the length of his sentence. Indeed, it would be impossible for a defendant to even know whether such a factor did influence his score, as the lack of transparency prevents a defendant from knowing the extent to which any given data (in public records or the questionnaire) has proved to be material. A defendant is therefore only given an opportunity to argue against a score in the absence of any real understanding of the basis for its calculation. Similar due process concerns because of lack of transparency also arise in parts of the SCS system.

Further challenges to equality before the law and due process safeguards can arise in some cases of automated decision-making due to what could be described as a 'reversal' of the burden of proof or lowering of the 'evidence threshold'. (On the importance of the burden of proof and 'evidence threshold' in the context of social welfare in the US, see Kaplow 2012: 738. In Australia, see, for example, Gray 2012: 13. In the context of the European Court of Human Rights, see Ambrus 2014. On due process implications of shifting the burden of proof in the US legal context, see McCauliff 1982: 1293; Petrou 1984: 822.). For example, in the robo-debt case, debt notices were issued for money that was not in fact owed by some welfare recipients, and the fact-finding burden for debt that previously rested on the Department was reversed, arguably contrary to the enabling legislation (Hanks 2017). While debts issued under this automated decision-making process can be challenged, it has been argued that the government failed its responsibility to ensure that it has established the existence of the debt before initiating the claim (Carney 2018).

Finally, the use of automated decision-making by governments poses a further challenge to the idea that all individuals irrespective of their status must have equal access to rights in the law, and that in accessing these rights 'like cases be treated alike'. This includes the notion that governments should not treat individuals differently due to their demographic group or an immutable trait. (People have particularly strongly objected to courts systematically imposing more severe sentences on defendants who are poor or uneducated or from a certain demographic group: see Kleck 1981: 783; Wacquant 2001: 401; Hsieh and Pugh 1993: 182.) Automated decision-making systems, such as COMPAS and Sesame Credit, can undermine this principle because they may: a) explicitly incorporate and rely on various static factors and/or immutable characteristics such as socio-economic status, employment and education, postal codes, age or gender; or b) take such matters into account indirectly, for example by 'learning' the relevance of variables that correlate with these. For example, in *Loomis*, the defendant had argued that the judge's consideration of the COMPAS score also violated his constitutional rights because COMPAS software used 'gendered assessments' (*Loomis* at 757) and, in turn, undermined his right to an individual sentence.

The greatest challenge to equality before the law comes from the fact that automation can infer rules from historical patterns and correlations. Even when variables, such as race, are not used in the learning process, a machine can still produce racially or otherwise biased assessments (Angwin, et al. 2016). This unequal treatment before the law results because many other factors can correlate with race, including publicly available information such as, for example, Facebook 'likes' which are not excluded from the machine learning process (see especially Kosinski, Stillwell, and Graepel 2013: 5802; finding that easily accessible digital records such as Facebook 'likes' can be used to automatically and accurately predict highly sensitive personal information, including sexuality and ethnicity). Further, the data from a pre–sentencing questionnaire (from which the COMPAS tool draws inferences) records the number of times and the first time a defendant has been 'stopped' by police. Given historical profiling practices of law enforcement in the United States, status as an African-American is likely to correlate with higher numbers and earlier ages in response to this question (O'Neil 2016: 25–26; 'so if early "involvement" with the police signals recidivism, poor people and racial minorities look far riskier'). Racial differentiation is thus built into the data from which correlations are deduced and inferences are drawn.

Unlike the risk assessment tool COMPAS, decisions in Swedish student welfare management system are made solely on factors that are legally relevant. The pre-programmed nature of the system ensures that those factors play a role in the decision precisely in the circumstances in which they are relevant. Decisions are made consistently with the law, with students treated equally under that law. In Chinese SCS, diversity of implementation means that equality before the law is affected differently. For example, decisions in the Rongcheng City model of the SCS system are made solely with reference to clearly defined categories of behaviour which leads to either a point deduction or addition – there is no room to consider any other factors in the pre-programmed system. In contrast, however, the Sesame Credit system in the SCS relies on variables that are irrelevant from a rule of law perspective, such as the rankings of an individual's social network contacts, which could lead to differential treatment in effect based on social status, sex or ethnic origin (see Kosinski, Stillwell, and Graepel 2013).

As our examples demonstrate, in understanding the benefits and challenges of automating government decisions, it is crucial to consider both the context of the decision and the type of system deployed. A system with pre-programmed rules can ensure that decisions are made based on factors recognised as legally relevant and hence avoid or minimise the risk of corruption or favouritism by officials. However, procedural rights and opportunities to check and rectify data on which the decision relies are crucial, as is ensuring that the logic of the system accurately reflects the law. As our case studies demonstrate, the challenges posed by systems based on rules inferred from data are different. Here, the role of humans is limited to setting parameters, selecting data (possibly biased due to flawed human collection practices), and deciding which variables to use as a basis for analysis. Unless the humans involved in these processes who have a deep understanding of the legal context in which a decision is made, systems may fail in practice to meet the standard of equality before the law. The

COMPAS system is an example of software that does not meet the needs of a fair criminal justice process – lack of transparency in a tool that relies on a large set of often legally irrelevant inputs prevents a defendant from having sufficient opportunity to participate in the court's findings on dangerousness, which is a crucial component of the ultimate decision.

This does not imply that machine learning techniques can never be used in government decision-making in ways that do comply with equality before the law. It can be used in the development of high-level policies, from traffic flow management to modelling interventions in the economy. Even at the level of decisions affecting individuals, machine learning is sometimes consistent with or even of benefit to equality before the law. For example, facial recognition, if designed to recognise the faces of diverse individuals accurately, could be used to identify individuals where that is an aspect of the system, and if programmed correctly may even overcome conscious and unconscious bias on the part of humans. While concerns about privacy and surveillance may counter its benefits, the use of machine learning in such a system can improve equality before the law by reducing arbitrariness.

D. Conclusion

Automation can improve government decision-making. The benefits include cost savings and greater speed, as well as a capacity to enhance the rule of law. Properly designed, implemented and supervised automation can help government decision-making better reflect the values of transparency and accountability, predictability and consistency, and equality before the law.

What is apparent, though, is that three of the four studies of automation considered in this chapter fail to live up to this ideal. In some cases, such as robo-debt, this failure results from poor design and implementation of the automated system. Indeed, one consistent theme is that human choices, and often error, at the design and implementation stage of automation can cause a system to fail to meet rule of law standards. A contrast is the Swedish student welfare system, which involves high levels of automation, but does not raise the same concerns. The Swedish model, which puts a strong emphasis on compliance with national legislation, officers' ethical codes, and publicity of the rules, demonstrates how a carefully designed system integrating automation with human responsibility can realise many benefits, while remaining sensitive to the values expressed in the rule of law.

It would nonetheless be a mistake to suggest that effective human design and implementation can ensure a particular automation technique will enhance or at least meet the minimum standards of the rule of law. It is clear from our study that even with active human engagement, some forms of technology raise intractable problems. This may be because the form of automation is inappropriate for its context. For example, machine learning offers many benefits, but some techniques or software products come at the price of transparency and accountability. This may be tolerable in particular circumstances, such as in the distribution of low-level welfare benefits (with appeal mechanisms), assisting with tasks such as optimising the traffic flow in a city, or conducting facial recognition for identification purposes. In such cases, testing and evaluation of accuracy and disparate impact may be sufficient from a rule of law perspective.

On the other hand, machine learning that cannot be rendered transparent and comprehensible may not be appropriate where it is used to make decisions that have greater effects upon the lives and liberties of individuals. It can also be inappropriate where a machine learning system may be influenced by criteria that ought not to be relevant, such as a person's race or even variables that have not traditionally been used to discriminate, such as the credit rating of one's friends. Such problems are exacerbated, as in the case of COMPAS, when the system operates according to undisclosed, proprietary algorithms. These problems would be compounded if COMPAS were used not only to assist judges, but to replace them.

From the perspective of the rule of law, these problems may become more acute over time. As technology develops, and machine learning becomes more sophisticated, forms of automation used by governments may increasingly become intelligible only to those with the highest level of technical expertise. The result may be government decision-making operating according to systems that are so complex that they are beyond the understanding of those affected by such decisions. This raises further questions about the capacity of voters in democratic systems to evaluate and so hold to account their governments, including in respect of compliance with rule of law values. Ignorance in the face of extreme complexity may enable officials to transfer blame to automated systems, whether or not this is deserved. The result may be an increasing tension between automation and the rule of law, even where humans design such systems in ways that seek to respect such values.

Ultimately, humans must evaluate each decision-making process and consider what forms of automation are useful, appropriate and consistent with the rule of law. The design, implementation and evaluation of any automated components, as well as the entire decision-making process including human elements, should be consistent with such values. It remains though to be seen whether these values can be fully integrated into automated decision-making and decision-support systems used by governments. Converting rule of law values into design specifications that can be understood by system designers, and enforced through regulation, professional standards, contracts, courts or other mechanisms, represents a formidable technical and legal challenge. This chapter highlights a number of common themes in this respect, including the need for an awareness of the link between tools and design, transparency and accountability, the need to consider consistency and predictability not only over time but also as between automated and human systems, the importance of embedding procedural due process rights, and the tension between deriving rules from historic data and equality before the law. Resolving these issues in the automation of government decisions will be critical for any nation that claims to uphold the basic ideals of the rule of law.

Bibliography

'State of Wisconsin v. Loomis' (2016), 881 N.W.2d 749.

Administrative Procedure Act, Pub L 79-404, 60 Stat 237, 5 USC §§ 551-559.

Ambrus, M. (2014), 'The European court of human rights and standards of proof: an evidential approach toward the margin of appreciation', in Gruszczynski, L. and Werner, W. (eds.), *Deference in international courts and tribunals: standard of review and margin of appreciation*, Oxford: Oxford University Press.

Angwin, J. (2018), 'Sample-COMPAS-Risk-Assessment-COMPAS-"CORE"', available at <https://www.documentcloud.org/documents/2702103-Sample-Risk-Assessment-COMPAS-CORE.html> (last accessed 16 August 2018).

Angwin, J., J. Larson, S. Mattu and L. Kirchner (2016), 'Machine bias', *ProPublica,* 23 May 2016, available at <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> (last accessed 16 August 2018).

Bingham, T.H. (2007), 'The rule of law', *CLJ* 66, p. 67.

Burrell, J. (2016), 'How the machine "thinks": understanding opacity in machine learning algorithms', Big Data & Society 3: 1, p. 1.

Carlson, A.M. (2017), 'The need for transparency in the age of predictive sentencing algorithms', *Iowa Law Review,* 103, p. 303.

Carney, T. (2018), 'The new digital future for welfare: debts without legal proofs or moral authority?', *UNSW Law Journal Forum*, May 2018, available at <http://www.unswlawjournal.unsw.edu.au/wp-content/uploads/2018/03/006-Carney.pdf> (last accessed 16 August 2018>.

CCJ/COSCA Criminal Justice Committee (2011), 'In support of the guiding principles on using risk and needs assessment information in the sentencing process', 3 August 2011, available at <http://ccj.ncsc.org/~/media/Microsites/Files/CCJ/Resolutions/08032011-Support-Guiding-Principles-Using-Risk-Needs-Assessment-Information-Sentencing-Process.ashx> (last accessed 15 August 2018).

Citron, D.K. and F. Pasquale (2014), 'The scored society: due process for automated predictions', *Washington University Law Review* 90, p. 1.

Commonwealth Ombudsman (2017), 'Centrelink's automated debt raising and recovery system: a report about the department of human services' online compliance intervention system for debt raising and recovery', Investigation Report, available at <https://www.ombudsman.gov.au/__data/assets/pdf_file/0022/43528/Report-Centrelinks-automated-debt-raising-and-recovery-system-April-2017.pdf> (last accessed 27 November 2018).

Creemers, R. (2018), 'China's social credit system: an evolving practice of control', available at

<https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3175792> (last accessed 16 August 2018).

CSN (n.d.), available at <https://www.csn.se/languages/english.html> (last accessed 6 November 2018).

Data Matching Program (Assistance and Tax) Act 1990 (Cth).

Diakopoulos, N. (2016), 'We need to know the algorithms the government uses to make important decisions about us', *The Conversation*, 24 May 2016, available at <http://theconversation.com/we-need-to-know-the-algorithms-the-government-uses-to-make-important-decisions-about-us-57869> (last accessed 16 August 2018).

Equivant (n.d.), 'Software for Justice', available at <http://www.equivant.com/> (last accessed 10 September 2018).

Gowder, P. (2016), *The rule of law in the real world*, Cambridge: Cambridge University Press.

Gray, A. (2012), 'Constitutionally protecting the presumption of innocence', *University of Tasmania Law Review*, 31: 1, p. 13.

Gunning, D. (n.d.), 'Explainable artificial intelligence (XAI)', Defense Advanced Research Projects Agency Project Information, available at <https://www.darpa.mil/program/explainable-artificial-intelligence> (last accessed 16 August 2018).

Hanks, P. (2017), 'Administrative law and welfare rights: a 40-year story from *Green v Daniels* to "robot debt recovery", *AIAL Forum,* 89, p. 1.

Hsieh, C. and M.D. Pugh (1993), 'Poverty, income inequality, and violent crime: a meta-analysis of recent aggregate data studies', *Criminal Justice Review*, 18, p. 182.

Hvistendahl, M. (2017), 'Inside China's vast new experiment in social ranking', *Wired*, 14 December 2017, available at <https://www.wired.com/story/age-of-social-credit/> (last accessed 10 September 2018).

International Congress of Jurists (1959), 'The rule of law in a free society', Report of the International Commission of Jurists, New Delhi.

Junwei, T. (2017), 荣成：建信用体系 创"示范城市 [Rongcheng: the making of a demonstration city for the social credit system], 新华社 [Xinhua News Agency], 13 July 2017, available at <http://xinhua-rss.zhongguowangshi.com/13701/6003014383535113117/2049163.html> (last accessed 10 September 2018).

Kaplow, L. (2012), 'Burden of proof', *Yale Law Journal*, 121, p. 738.

Kehl, D.L., P. Guo and S.A. Kessler, (2017), 'Algorithms in the criminal justice system: assessing the use of risk assessments in sentencing', *Berkman Klein Center for Internet & Society, Harvard Law School*, available at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:33746041> (last accessed 16 August 2018).

Kennedy, R. (2017), 'Algorithms and the rule of law', *Legal Information Management*, 17, p. 170.

Kleck, G. (1981), 'Racial discrimination in criminal sentencing: a critical evaluation of the evidence with additional evidence on the death penalty',*American Sociological Review*, 46, p. 783.

Kosinski, M., D. Stillwell and T. Graepel (2013), 'Private traits and attributes are predictable from digital records of human behaviour', *Proceedings of the National Academy of Sciences of the United States of America*, 110, p. 5802.

Launchbury, J. (2017), 'A DARPA perspective on artificial intelligence', DAPRAtv, YouTube, available at <https://www.youtube.com/watch?v=-O01G3tSYpU> (last accessed 20 August 2018).

Mathiesen, T. (1998), 'Selective incapacitation revisited', *Law and Human Behaviour*, 22, 455.

McCauliff, C.M.A. (1982), 'Burdens of proof: degrees of belief, quanta of evidence, or constitutional guarantees', *Vanderbilt Law Review*, 35, p. 1293.

Mikhaylov, S.J., M. Esteve and A. Campion (2018), 'Artificial intelligence for the public sector: opportunities and challenges of cross-sector collaboration', *Philosophical Transactions of the Royal Society A*, 376: 2128, available at <https://doi.org/10.1098/rsta.2017.0357> (last accessed 10 September 2018).

Mistreanu, S. (2018), 'Life inside china's social credit laboratory', *Foreign Policy* (online), 3 April 2018, available at <https://foreignpolicy.com/2018/04/03/life-inside-

chinas-social-credit-laboratory/> (last accessed 10 September 2018).

Nisbett, R.E. and T.D. Wilson (1977), 'Telling more than we can know: verbal reports on mental processes', *Psychological Review*, 84, p. 231.

Note (1982), 'Selective incapacitation: reducing crime through predictions of recidivism', *Harvard Law Review*, 96, p. 511.

Noto La Diega, G. (2018), 'Against the dehumanisation of decision-making: algorithmic decisions at the crossroads of intellectual property, Data Protection and Freedom of Information', *JIPITEC* 9, p. 3.

O'Neil, C. (2016), *Weapons of math destruction: how big data increases inequality and threatens democracy*, New York: Broadway Books.

OKS (2014), 'Dnr: 2014-03172', 13 October 2014, available at <https://oks.se/wp-content/uploads/2016/03/2014-03172.pdf> (last accessed 6 November 2018).

OKS (n.d.), 'Avgöranden (rulings)', available at <https://oks.se/avgoranden/> (last accessed 6 November 2018).

Perry, M. (2017), 'iDecide: administrative decision-making in the digital world', 91 *Australian Law Journal,* 91, p. 29.

Petrou, P. (1984), 'Due process implications of shifting the burden of proof in forfeiture proceedings arising out of illegal drug transactions', *Duke Law Journal* 822.

Schuerman J.R., E. Mullen, M. Stagner and P. Johnson (1989), 'First generation expert systems in social welfare', *Computers in Human Services,* 4, 111.

Senate Community Affairs References Committee (2017), 'Design, scope, cost-benefit analysis, contracts awarded and implementation associated with the better management of the social welfare system initiative', Parliament of Australia.

Shanghai Development and Reform Commission (2017), '上海市社会信用条例 [Shanghai social credit regulations]', 29 June 2017, available at <http://www.shdrc.gov.cn/gk/xxgkml/zcwj/zgjjl/27789.htm> (last accessed 10 September 2018).

State Council (2014), 'Planning outline for the construction of a social credit system (2014–2020)', 14 June 2014, in Creemers, R. (ed.) (2015), *China Copyright and Media*, 25 April, <https://chinacopyrightandmedia.wordpress.com/2014/06/14/planning-outline-for-the-construction-of-a-social-credit-system-2014-2020/> (last accessed 16 August 2018).

Susskind, R.E. (1987), *Expert systems in law: a jurisprudential inquiry*, Oxford: Clarendon Press.

Sutcliffe, J. (1989), 'Welfare benefits adviser: a local government expert system application', *Computer Law & Security Review,* 4: 6, p. 22.

Swedish Tax Agency (2016), 'Taxes in Sweden: an English summary of tax statistical yearbook of Sweden', available at <https://www.skatteverket.se/download/18.361dc8c15

312eff6fd1f7cd/1467206001885/taxes-in-sweden-skv104-utgava16.pdf> (last accessed 10 September 2018).

Tamanaha, B.Z. (2004), *On the rule of law: history, politics, theory*, Cambridge: Cambridge University Press.

Tyree, A. (1989), *Expert systems in law*, Sydney: Prentice Hall.

Wacquant, L. (2001), 'The penalisation of poverty and the rise of neo-liberalism', *European Journal on Criminal Policy and Research*, 9, p. 401.

Wihlborg, E., H. Larsson and K. Hedström (2016), '"The computer says no!" A case study on automated decision-making in public authorities', 49th Hawaii International Conference on System Sciences.

Zhong, R. and P. Mozur (2018), 'Tech giants feel the squeeze as Xi Jinping tightens his grip', *New York Times* (online), 2 May 2018, available at <https://www.nytimes.com/2018/05/02/technology/china-xi-jinping-technology-innovation.html> (last accessed 10 September 2018).