

Governing algorithmic decision-making in government

The role of the Senate



Memorandum for the AI working group

In both society and in politics, more and more attention is given to both the possibilities and the concerns related to the use of artificial intelligence (AI) in government. AI can reveal connections and insights that are difficult for humans to notice. In the context of Justice and Security (JenV), it offers opportunities in the field of detection and intelligence. However, it takes effort to develop and implement such AI applications in a responsible way. Therefore, to explore the Senate's options to contribute to responsible development and implementation of AI decision-making within the government, the JenV committee has established an AI working group .

The Rathenau Instituut's task is to contribute to the societal dialogue and the formation of political judgements on the impact of scientific and technological developments on society. At the request and with input from the Senate, the Rathenau Instituut wrote this memorandum, intended to inform the research of the AI working group. Different

commissions are involved in the safeguarding of fundamental rights and democratic control in the government's use of AI. The memorandum serves as preparation for a number of in-depth sessions and a senate-wide expert meeting on this complex and committee-transcending task.

The memorandum consists of:

1. A brief summary of some of the concerns in society, science and government about algorithmic government decision-making.
2. An overview of how the government has addressed these concerns in the period 2018-2020.
3. A list of challenges for the Senate with four points of discussion for the expert meeting.

1. Concerns about algorithmic decision-making by the government

The use of 'smart' systems to inform and take governmental decisions fits within a broader digitisation trend that has taken place in recent decades. For example, the government has long used digital systems to grant benefits, issue licences or automatically collect traffic fines. These are often systems based on so-called 'if, then' instructions, also known as rule-based AI.¹ Such a system can consist of simple rules (e.g. a clear decision tree), but can also be very complex due to many variables and rules.

For about five years now, the government has been using so-called 'smart' systems that are based on machine learning. These systems discover connections in large quantities of data and, on the basis of that data, calculate, for example, the probability that someone has committed fraud (System Risk Indication, SyRI), or that certain forms of crime occur in a certain neighbourhood (predictive policing). They are learning systems, in which ICT professionals teach the system to recognise patterns using training data.²

Both types of systems can be used in government decision-making, both for support and for automatic decision-making. It is not always easy to determine whether the government uses (self)learning systems or 'just intensive data use and linking'. When it comes to 'data-intensive' applications, the expectation is that the next step will be the use of (self)learning algorithms (Wolswinkel, 2020a; TNO, 2019). In practice, we

¹ A rule-based system does not learn from itself, but does exhibit intelligent behaviour by analysing the environment and taking actions - with some degree of autonomy - to achieve specific goals (see also the European Commission's definition of AI, EC 2019). Much attention around AI now focuses on self-learning systems (see footnote 2), with some no longer considering rule-based systems to be true AI (Rathenau Institute, 2019).

² Learning algorithms can be trained in different ways. In supervised learning, the algorithm is trained on the basis of labelled examples, for example a picture of a cat is labelled 'cat', a picture without a cat is labelled 'no cat'. Programmers teach the algorithm to deal with the examples provided. In unsupervised learning, the training data are not labelled. The algorithm is then programmed in such a way that after some time it can independently recognise structures or patterns in the input data. (Kulk & van Deursen, 2020).

already see hybrid forms emerging; decision trees can be combined with algorithms that are *trained*.³

Box 1 How 'smart' is AI?

As of yet, most machine learning systems are good at one specific task. In some areas, such as image recognition, the software is making rapid progress. In other areas, such as making predictions about complex social outcomes related to fraud or job suitability, progress is less straightforward.

This is related to the 'problem space' (Schermer, et al., 2020). In a chess computer, for example, the environment in which an AI system operates is clear. There is a game board and there are unambiguous rules within which specific objectives must be achieved. In facial recognition software, too, the problem is fairly well defined. Because systems can be fed with a lot of photos and videos of faces, they are increasingly successful in identifying faces accurately.

But in more complex and dynamic environments, there are many more variables. There are numerous factors that, for example, a self-driving car must take into account. It is also difficult to clearly define priorities or desirable strategies; if problems arise, will the car primarily protect the occupant or other road users?

When profiling human behaviour and predicting social outcomes, the 'problem space' is even larger. AI systems are being developed to select who is the most suitable job applicant or when and where someone is going to show criminal behaviour. But it is not clear which data is relevant and which dimensions can be captured in data. When the problem space can no longer be "written out", the chance of unforeseen effects and errors increases.

In this paper, we focus on issues concerning the rule of law that arise because systems (a) sometimes operate in a less than transparent, or even opaque manner, (b) often leave little room for individual circumstances, and (c) can violate fundamental rights, such as privacy and non-discrimination. We briefly address each of these points.

1.1 Algorithmic systems are not always transparent

The algorithms and choices of rule-based AI systems are programmed by humans, and, therefore, humans should be able to understand and know how these systems operate. However, in the case of system development by third parties - when the software is protected by intellectual property - it can be difficult for the government and parliament to guarantee transparency about the choices made by the software. There are also other cases in which it can be difficult to determine how decisions of rule-based systems are made.

3 See, for example, https://www.dsp-groep.nl/wp-content/uploads/11onbetaal_Onderzoek_gebruik_betaalprofielen_DSP-Decisio_2019.pdf

Box 2 Value judgements of ICT professionals

During system design, developers and data analysts are constantly making choices and value judgements - it is impossible not to (Wieringa 2020; RvS 2018). The translation of laws and regulations into algorithms often requires the prioritisation and qualitative expression of value judgements. If value judgements and political choices also take place during the legislative process, these choices are tested and weighed by parliament - a democratic process takes place. This is not the case when IT professionals have to interpret the law and make value judgements. Parliament has no supervision or control over this.

Sometimes, legal texts are worked out in such detail that programmers can copy the descriptions and principles almost seamlessly as code. For example, in the *Wet beslagvrije voet* (section 475da), precise amounts and calculation rules are described for each category - a single person, a single parent, married people without or with children (Passchier, 2021).

Sometimes, however, laws and regulations leave a lot of room for interpretation at the executive level and the same situation can be weighed differently (Van Eck et al., 2018). In order to decide whether to withdraw a licence, for example, it must be assessed which defects are serious enough. When criteria are multi-interpretable or someone can meet a criterion to different degrees, value judgements must be made when programming an AI system. ICT professionals who do not formally have decision-making authority then become *de facto* policymakers.

There are several ways to increase transparency and democratic control. One option is involving ICT professionals in the drafting of legal texts. This can help prevent confusion in the translation of laws into AI systems at a later stage.⁴ Another option is organising mandatory impact assessments or audits, for example by organisations themselves, to help enforce that ICT professionals reflect on the choices they make and that these are recorded in such a way that, afterwards, they are transparent and verifiable by citizens, parliament, the supervisor and/or the judge.

The choices of machine learning systems are, at a fundamental level, more difficult to follow and to explain than the choices of rule-based AI systems. There are various machine learning techniques in which it is in fact not the ICT professionals, but the systems themselves, that determine what the relevant criteria or patterns are. This process relies heavily on probability (statistics) and is accompanied by an enormous 'data hunger'.⁵ This does not mean, however, that people are no longer involved in the decision-making process of systems. Software developers teach the system to recognise patterns with the help of training data, so that the system will eventually be able to analyse data sets that humans are not able to process. The AI system is

4 Science and government are looking for ways to achieve "machine enforceable legislation". An example is the *Calculus-FLINT* project, which looks at how design choices can already be taken into account when drafting legislation. Parliament can review these better designed laws, increasing the democratic legitimacy of AI systems.

5 The availability of large amounts of data is crucial for training AI systems and making them better (see footnote 2).

dependent on the data it receives as input. If the data contains errors, the AI system will not generate sensible results. In addition, problems may arise when correlations are interpreted as causal relationships. In machine learning systems, designers can look inside the system and adjust the programming code. But this often provides only limited insight into the way in which the system handles all kinds of criteria and data.⁶

This means that it is possible that choices made by ICT professionals and/or AI systems during software development will evade political debate and democratic control. If it remains unclear how (and by whom) value judgements are structurally converted into numbers, or how it is determined when a system is 'good enough' and what acceptable error margins are (Wieringa, 2020), this has fundamental consequences for the democratic legitimacy of AI decisions.

1.2 AI systems leave little room to take individual circumstances into account

While ICT professionals gain more influence to translate and implement policy (as they see fit), civil servants' autonomy is diminishing (see Bovens & Zouridis 2002; Zouridis 2000). Twenty years ago, public administration experts already described the transformation that took place within the government through the application of ICT. Instead of a civil servant deciding on an individual case, there is a system of algorithms that comes to a decision (Bovens & Zouridis 2002). Decisions are becoming more standardised and more and more scenarios and solutions are already anticipated at system level. Various laws also leave little room for discretion, such as the legislation that played a role in the 'toeslagenaffaire' (benefits scandal). The flexibility to formulate appropriate solutions for individual cases can be lost.

Box 3 Heartbreaking cases

One mistake does not make you a fraudster. Nevertheless, tens of thousands of parents got into serious financial problems because of the toeslagenaffaire (benefits scandal). The report 'Ongekend Onrecht' (Unprecedented Injustice) calls for critical reflection on the functioning of the entire government, including the use of algorithmic government decision-making. A specific AI system must be able to be adapted quickly and effectively when things go wrong, because in some AI applications, small errors in a database can have major consequences for citizens.

An error in the Automatic Number Plate Recognition (ANPR) file, for example, can result in someone being stopped by the police more than once. This error may have occurred in the delivery of information by the Centraal Bureau Rijvaardigheid (CBR), the Dienst Wegverkeer (RDW) or the municipality. The Nationale Ombudsman ruled that the police are nevertheless responsible for a reliable ANPR file and must ensure that errors are rectified (Nationale Ombudsman, 2015).

6 Science is working on methods to improve this, see the Hybrid Intelligence project in the NWO Gravity Programme.

Unfortunately, it regularly happens that it is (almost) impossible to quickly trace an error and fix it with the push of a button in all interconnected systems (Widlak and Peters 2018; Van Eck et al., 2018).⁷ In recent years, the Nationale Ombudsman and the Kafkabrigade Foundation have documented harrowing cases of citizens who become trapped in the digital government bureaucracy.

When systems take decisions on the basis of incorrect, outdated or incomplete data, the responsibility soon lies with the citizen to spot errors and start an objection procedure (WRR, 2017). Also, there is often no human point of contact and citizens are sent from pillar to post (TCU, 2021).

It is important to ensure that it remains possible to intervene when systems crash or when unforeseen effects occur that have an undesired impact on citizens. In the case of AI systems, the design can determine the desired degree of human involvement and flexibility.⁸ One possibility is mandatory transparency about the origin of the data used, so that it can be corrected or contested.

However, it is not just about the technology. It is crucial to build in an option to further explain a specific situation to a human, who can then also make a judgment and deviate from the AI decision. When there is not enough room for people to correct the system or (re)assess specific outcomes themselves, the human touch and meaningful human intervention is compromised.

1.3 AI systems may violate fundamental rights

In recent years, several studies have shown that algorithmic government decision-making poses risks to the protection of fundamental rights (Rathenau Instituut 2017a; Vetzo et al. 2018; WRR 2016). Problems can arise in the implementation of systems. For example, decent treatment and the right to human contact are under pressure when citizens are confronted with 'computer says no' scenarios (see also Box 3). But things can also go wrong when choosing to build a system and during the development phase of the system, resulting in, for example, undesirable discrimination (see box 4). In recent years, the right to respect for privacy, the right to equal treatment and the right to a fair trial have been terribly compromised in various cases.

Box 4 Types of discrimination

7 It is often an option to make a note of an error, but then problems can still arise due to chain automation in the government. In IT systems, this sometimes makes it virtually impossible to correct errors retrospectively, whereas this is an important principle in law.

8 A distinction is often made between human-in-the-loop (human gives permission; AI system as a tool), human-on-the-loop (human can intervene; AI system can be stopped) and human-out-of-the-loop (no human supervision; automatic decision-making). But also at the level of digital forms, for example, much can be gained by taking citizens' wishes into account and working on user-friendliness (Nationale Ombudsman, 2019).

There are many ways in which biases can creep into an AI system and, as a result, the AI system makes systematic discriminatory decisions (Barocas and Selbst, 2016; Zuiderveen-Borgesius 2019). Problems can occur due to the difficulty of defining the system's mission and relevant features or variables (see also Box 1). But things can also go wrong at the level of training data. These may, for example, be non-representative or already contain human bias.

For example, the ICT industry has traditionally employed more men than women. Such a historical fact may cause an AI system to unintentionally reproduce an 'old' bias when selecting applicants.⁹ Choosing an Ivy League degree as a selection criterion may also be problematic. After all, these universities are not (always) inclusive in their selection process either. Moreover, according to scientists, diplomas are not the best indicator of future success.¹⁰ But even a criterion such as 'never be late for work' can have a discriminatory effect, as certain groups are more at risk of delays due to their dependence on public transport or more affordable housing outside the city.

Finally, certain characteristics may also be coded in other data. If information about gender or skin colour is not included in the training data, such characteristics may still be considered in the decision of an AI system because correlations are known, for example, between postcodes and ethnicity. The way in which training data is collected can also result in a biased dataset. When the police focus more on certain ethnic groups or certain neighbourhoods, this does not result in a fair sample. It is then not surprising that those groups or neighbourhoods are systematically over-represented in police records.

Various AI systems are the subject of social and political debate due to their possible violation of fundamental rights. Some systems have been banned by the courts, such as the SyRI system that uses algorithms to profile citizens in order to detect fraud. According to the judge, the social interest and the infringement of civil rights were not properly balanced. The introduction of AI systems therefore requires a careful fundamental rights review.

2. The government's toolkit

Given the above-mentioned concerns, many consider better democratic accountability for the development and use of algorithmic decision-making to be necessary (CAHAI 2020; EC 2019). The government has also largely acknowledged the concerns described above. Over the past three years, the government has considered in various studies and parliamentary letters whether the existing legal framework and the associated oversight system - the government's 'toolkit' - are sufficiently future-proof and has investigated what is needed to mitigate concerns about accountability, human

9 This happened to Amazon when they were developing a system for automatically selecting job applicants. Even when the developers became aware of this particular bias in the system, they could not get it out, and the company was forced to stop developing the system.

10 See: Rathenau Instituut (2020a).

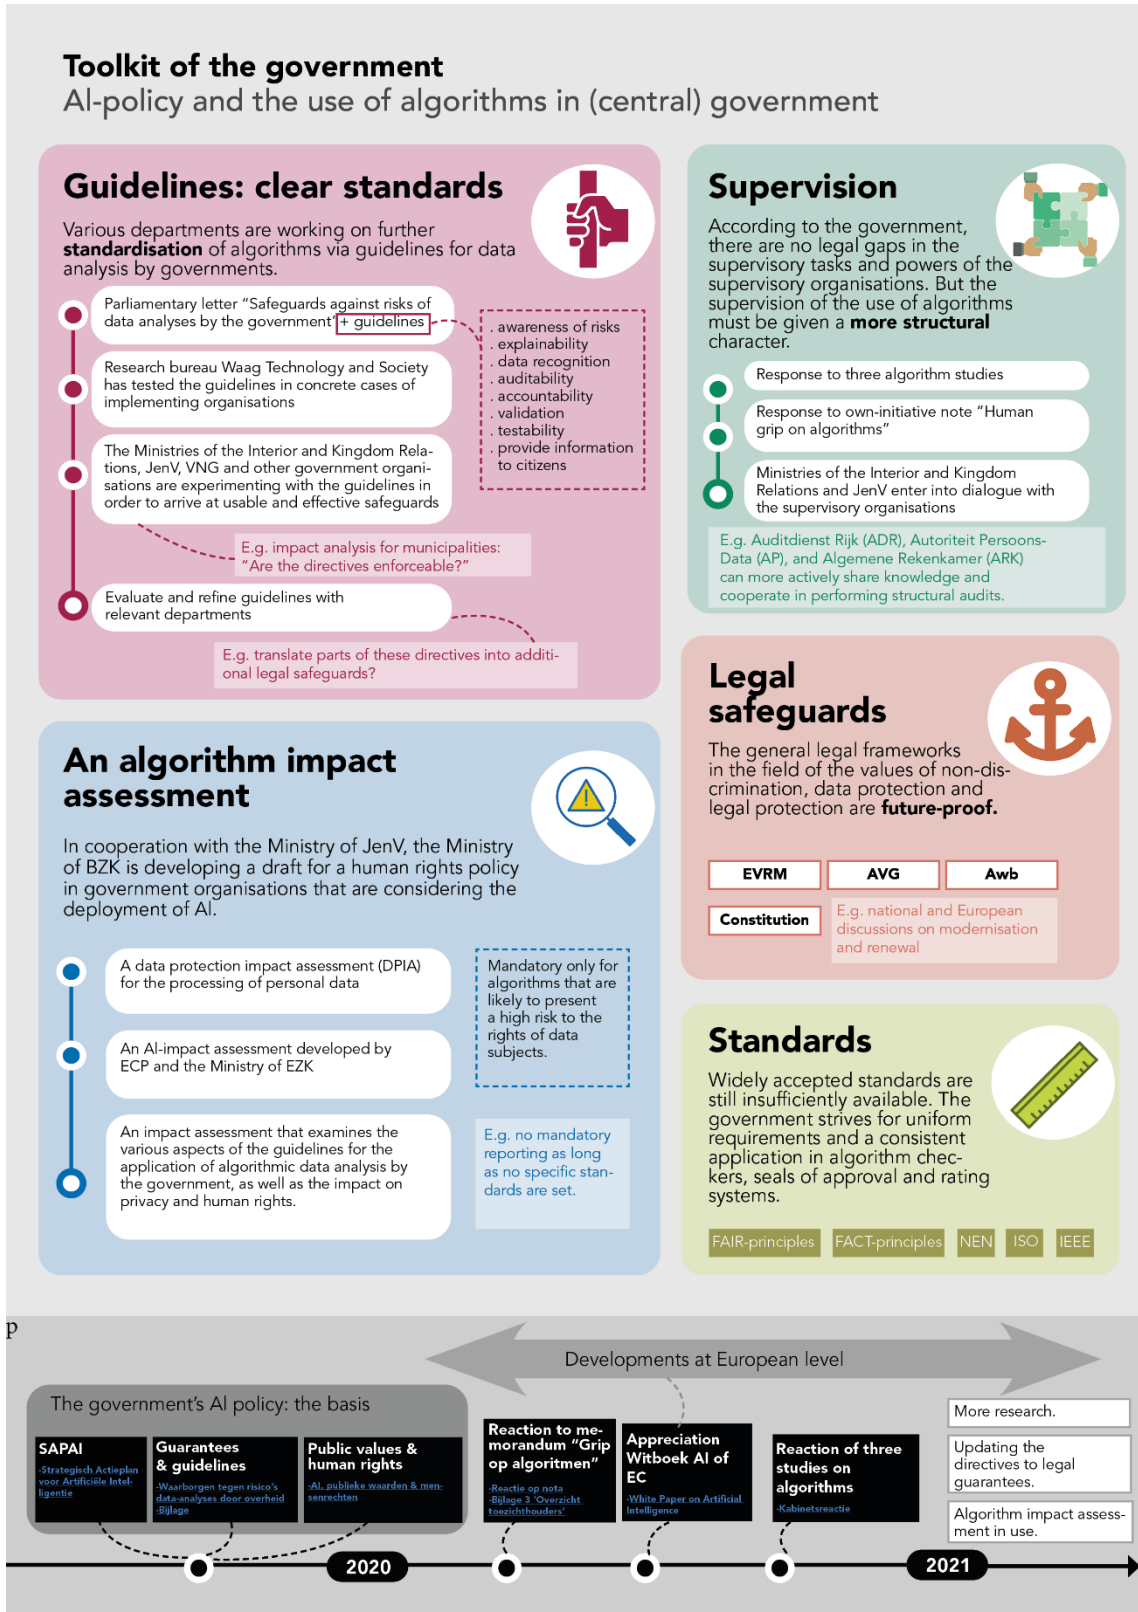
standards and human rights. We present the main actions taken by the government in Figure 1.

Important frameworks are the Constitution, the Algemene wet bestuursrecht (General Administrative Law Act) or Awb, and the Algemene Verordening Gegevensbescherming (General Data Protection Regulation) or AVG. The overall picture that emerges from the documents is that the legal framework and supervision are largely in order, but that there are still gaps that need to be filled (Parliamentary Papers II 2019-2020, 26643, no. 726). In addition, the government considers it necessary to provide more structural supervision of algorithmic decision-making by the Algemene Rekenkamer (Court of Audit), the Autoriteit Persoonsgegevens (Personal Data Authority) and the Auditdienst Rijk (National Audit Office) (Parliamentary Papers II 2019-2020, 35212, no. 3).¹¹

In particular, the government wants to focus on the question of whether the current information and transparency obligations of automated decision-making and profiling are sufficient (Parliamentary Papers 2019-2020, 26643, no. 641). For example, Article 22 of the AVG requires that fully automated decision-making must comply with various safeguards. But the AVG does not make clear exactly what safeguards apply to *partially* automated decision-making, while much algorithmic decision-making by the government falls within this category.

¹¹ The government has appointed an official working group to draw up a research agenda in the field of standardisation and supervision of algorithms.

Figure 1: Overview government policy for algorithmic decision-making



Another problem arises with so-called 'area-specific' data analysis. Such risk analyses are not focused on individuals, but on neighbourhoods or streets. Therefore, they do not fall under the AVG definition of profiling.¹² The government indicates that some of the risks of profiling, such as the chance of prejudice, may also occur with these data analyses.

The government is now looking to the European Commission for further regulation of AI systems (Kamerstukken II 2019-2020, 26643, nr. 726).¹³ It is expected that 'high-risk' AI applications will be subject to further regulatory requirements (EC, 2020; Kamerstukken II 2019-2020, 26643 / 32761, nr. 697).

In the meantime, the government is working on the development of guidelines, design principles and impact assessments to help system designers and implementing bodies to estimate and prevent possible human rights violations in advance, and to better justify the choices made during system design (Parliamentary Papers 2019-2020, 26643, no. 641; Parliamentary Papers II 2019-2020, 26643, no. 726). The government opts for a 'learning approach' and is funding, for example, the 'Calculemus-FLINT' project, in which a system is being worked out to transparently translate law into algorithm rules.

There are also considerations with regard to compliance with the Algemene wet bestuursrecht (General Administrative Law Act) or Awb. The Awb regulates, among other things, that citizens can submit their complaints to the Ombudsman and that both the Algemene Rekenkamer (General Court of Auditors) and the Auditdienst Rijk (Audit Department of the Central Government) audit the government. The principles of good governance, laid down in the Awb, also apply to algorithmic decision-making. The principles of reasoning and due care stipulate that the government must ensure that the choices made by algorithms are verifiable.¹⁴

3. Discussion points for the Senate

In addition to the efforts of the government, there are also ways for the Senate to strengthen its role with respect to algorithmic decision-making within the government. In order to inform an expert meeting on this subject and on the basis of published studies, advisory reports and ministerial letters, below, the Rathenau Instituut proposes four points of discussion.

12 For example, the police use area-based analysis to gather information on where certain forms of crime occur in a city, and to determine the required police deployment based on this information.

13 More specifically to the European Commission's risk approach with regards to AI, where it is expected that 'high-risk' AI applications will be subject to further regulatory requirements (EC, 2020), and the kabinetssapportage Witboek on Artificial Intelligence (Kamerstukken II 2019-2020, 26643 / 32761, nr. 697).

14 This transparency is not absolute. That is, insight into the choices made by the system do not mean that disclosure of the full source code is necessary (see Parliamentary Papers 2018 - 2019, 26 643, no. 570). For a further discussion of the future-proofing of administrative law in the context of self-learning systems, see, for example, Wolswinkel (2020a; 2020b).

The core task of the Senate is to assess the quality of legislation in a formal sense. Therefore, it has to be alert to various legislative proposals and amendments that come along and critically test them on essential criteria, such as legality, feasibility and enforceability. However, not every AI system is introduced by means of a law or legislative amendment, because some AI systems are regulated by existing legal frameworks. Therefore, the first discussion point concerns the various ways of anchoring algorithmic decision-making more firmly in law. Some legislative proposals and amendments may contribute to the proper development and use of all AI systems. Alertness to changes in such overarching laws and regulations is therefore crucial.

It also happens that a new law or amendment of the law immediately creates room to develop and apply certain digital systems. This was the case, for instance, with SyRI, which is linked to the Work and Income Implementation Structure Act (SUWI). When it is up to the Senate to assess such a bill, *three aspects* play an important role for a careful assessment:

- the proportionality of the proposed legislation,
- the way human intervention and supervision is embedded, and
- the extent to which parliament can monitor the development of the AI system after the adoption of the law.

3.1 How can we embed AI systems more strongly in the law?

We mentioned above that not every AI system is introduced by means of a law or an amendment to the law. For example, the Criminaliteits Anticipatie Systeem (Crime Anticipation System), or CAS, falls under the existing wet Politiegegevens (Police Data Act). As a result, CAS has never been discussed in the Dutch Senate, while there is now controversy about the invasion of privacy of citizens by the system, about alleged discriminatory effects and about the effectiveness of the predictions of CAS. The question arises whether systems with such an impact on fundamental rights require more parliamentary oversight.

One way to increase parliamentary control is through further general legal framing. Existing laws, after all, often already oversee the use of AI systems by the government, or set requirements for the government's actions. The Uitvoeringswet AVG and the Wet Politiegegevens (Police Data Act), for example, supervise data processing by various government departments, while the Algemene Bestuurswet (General Administrative Law Act) imposes requirements on the decisions made by the government - even if they are supported by algorithms. These laws do not regulate the arrival of a specific system, but they do indicate what the government or implementing bodies may or can do. Are adjustments or changes needed in these laws?

- ❖ *During the expert meeting the Senate can discuss various ways of anchoring AI systems more firmly in the law. This could include possible updates or enhancements to general laws that the Senate has to review, such as the*

Constitution, the General Administrative Law Act (Awb), the Act on Civil Procedure (WoB), the General Data Protection Act (AVG) and the Police Data Act. What, for example, are the advantages and disadvantages of the advice of scientists and the Council of State to give software code a higher legal status¹⁵?

3.2 How to support the Senate in the proportionality test of AI-related legislation

It is the task of the Senate to apply the principles of proportionality and subsidiarity in the context of protecting fundamental rights. In doing so, it is important to consider whether a possible infringement of these rights by AI-related legislation is justified, given the expected social benefit of the proposed system or type of system. It is also important to consider whether there are alternatives available that could achieve the same result, but which would be less or not at all harmful to fundamental rights.

In order to test this carefully, information is needed on the effectiveness and development of a proposed AI system. This can be difficult in advance, see discussion point 4 (section 3.4). It is also important to find out how exactly the system will be anchored in law. This will partly consist of a critical study of the advice of the Council of State and the supervisors involved, and of the government's response to these. The Senate also regularly organises expert meetings to prepare for certain legislative processes. Insights from human rights impact assessments, made by an implementing organisation, can also help the Senate members in their assessment.

The decision-making process of the Senate itself is divided into fixed committees, but one AI system can fall under several committees. A good proportionality check, therefore, requires a more integral treatment of a bill, involving several committees. The Senate can explore whether, and in what way, it can strengthen its grip on this decision-making process. For example, as a result of the Tijdelijke onderzoekscommissie Digitale toekomst, the House of Representatives has decided to establish its own standing committee on digital affairs, which can support the standing committees in decisions on AI (House of Representatives, 2020).¹⁶

- ❖ *During the expert meeting members can discuss how the Senate can further develop its knowledge and expertise. What information is needed: should reports such as the human rights impact assessments be sent to parliament on a mandatory basis? Is an integral treatment of legislative proposals necessary?*

¹⁵ The decision rules in software could, for example, be given the status of policy rules (in the sense of the Awb) or of an Order in Council (Algemene Maatregel van Bestuur, AMvB) (Van Eck et al. 2018; RvS, 2018).

¹⁶ This committee can monitor a member of the Cabinet responsible for digitisation, explore and put on the agenda future cross-committee developments, inform other standing committees and take the lead in dealing with cross-committee digitisation issues, and act as a contact point for digitisation issues for both the Cabinet and social groups, industry, science and others (House of Representatives 2020).

*Is it desirable to have a permanent committee on Digital Affairs in the Senate?
Is cooperation with the House of Representatives an option?¹⁷*

3.3 How can the Senate guarantee effective supervision and room for taking into account individual circumstances?

In addition to lawfulness, the Senate also tests for feasibility and enforceability. The degree of room for assessment or even policy freedom at the executive level is an important aspect here. The aim is to prevent potential confusion in the translation of the law into practice as much as possible. After all, it is undesirable to give decision-making authority to ICT professionals or even AI systems without democratic accountability. However, it is crucial to leave room to tailor to people's individual circumstances. Officials must be able to formulate an appropriate solution for an exceptional case. To prevent citizens from being confronted with 'computer says no' scenarios, there should be some flexibility at system level. It is therefore important to check whether new systems within official organisations offer sufficient scope to do this.

In this context it is important that the Senate looks at supervision. The various supervisory bodies, such as the Autoriteit Persoonsgegevens (Personal Data Authority) and the Algemene Rekenkamer (Court of Audit), have, on the basis of their functions and powers, possibilities for investigating the extent to which algorithmic decision-making in the government takes place in a responsible manner. Yet in practice this is easier said than done; the question arises to what extent this is sometimes only a reality on paper. All too often, supervision does not appear adequate to detect poorly functioning automated decision-making in time and, where necessary, to sound the alarm. It is therefore important that the Senate is able to carefully assess whether a legislative proposal offers enough scope for human intervention and checks whether effective supervision is possible. Here, too, the Senate must build up and use internal and external expertise.

- ❖ *During the expert meeting members can discuss how the Senate can get control of AI systems used in government decision-making: is it clear in advance which responsibilities various parties will have? In what way can the leeway of civil servants with regard to automated decision-making be laid down by law? And are supervisors equipped to carry out their task effectively ?*

3.4 What level of parliamentary involvement is desirable after the introduction of AI-related bills?

The review of an AI-related bill usually takes place before the development and actual introduction of that AI system. We mentioned above that during the system design phase, legislative provisions must be translated into algorithms. Parliament has no oversight or control over interpretations and value judgements that system designers

17 Various examples of parliamentary cooperation can be found abroad. In the United Kingdom, there are joint committees between both Houses, and enquiry committees in the House of Lords influence political discussions in the House of Commons. In order to get a better grip on digitisation issues, the House of Representatives has commissioned an inventory of the working methods of foreign parliaments (Rathenau Instituut 2020b).

are forced to make. The discretionary space of ICT professionals is increasing, but is still not standardised. In other words, little democratic control is currently being established around the implementation of the law in (future) AI systems.

In addition, little is usually known in advance about the effectiveness of smart systems. So far, for example, there appears to be little evidence that SyRI leads to better detection of fraud, or that predictive policing works better than traditional detection methods (Meijer and Wessels, 2019). Often, more information on the effectiveness will have to be gathered during the lifetime of the system, for example through a pilot or experiment. But if the Senate does not request this in advance, they will not have this information at a later stage, nor the possibility to weigh and test proportionality again.

Moreover, the circumstances under which a system started - and the political choices and considerations that were initially made - may change over time. AI software is so complex that it almost always has unforeseen effects that may have an undesired impact on citizens and businesses. Also, implementing agencies can push the limits of their mandate and, in retrospect, the mandate given may prove to be too broad.

For all these reasons, it may be desirable to reconsider the decision that the Senate initially made.

- ❖ *During the expert meeting members can explore ways in which the Senate can keep up to date and (re)evaluate relevant legislation at some point after its introduction. This could include discussing the advantages and disadvantages of temporarily valid legislation and the inclusion of sunset clauses.*

Bibliography

Barocas, S. en A. Selbst (2016). "Big Data's Disparate Impact." *Californian Law Review* 104, 671.

Bovens, M. en S. Zouridis (2002). "Van street-level bureaucratie naar systeem-level bureaucratie Over ICT, ambtelijke discretie en de democratische rechtsstaat." *Nederlands Juristenblad* 77(2): 65-74.

Europese Commissie (2020). *White paper on Artificial Intelligence: a European approach to excellence and trust*. https://ec.europa.eu/info/publications/white-paper-artificial-intelligence-european-approach-excellence-and-trust_en

Europese Commissie (2019). *A definition of AI: main capabilities and disciplines*. High-Level Expert Group on Artificial Intelligence. <https://ec.europa.eu/digital-single-market/en/news/definition-artificial-intelligence-main-capabilities-and-scientific-disciplines>

Kamerstukken II 2019-2020. 26643, nr. 726.

Kamerstukken II 2019-2020. 35212, nr. 3.

Kamerstukken II 2019-2020. 26643, nr. 641.

Kamerstukken II 2019-2020. 26643 / 32761, nr. 697.

Kulk, S. en M. van Deursen (2020). *Juridische aspecten van algoritmen die besluiten nemen – een verkennend onderzoek*. Den Haag: WODC 2020.

Meijer, A. en M. Wessels (2019). "Predictive Policing: Review of Benefits and Drawbacks." *International Journal of Public Administration* 42(12): 1031-1039.

Nationale Ombudsman, Jaarverslag 2019: 'Regel regie!'.

Nationale Ombudsman, Jaarverslag 2015: 'Burgerperspectief: een manier van kijken'.

Passchier, R. (2021). *Artificiële intelligentie en de rechtsstaat. Over verschuivende overheidsmacht, Big Tech en de noodzaak van constitutioneel onderzoek*. Den Haag: Boom Juridisch.

Raad van State (2018). *Ongevraagd advies over de effecten van de digitalisering voor de rechtsstatelijke verhoudingen*. <https://www.raadvanstate.nl/@112661/w04-18-0230/>

Rathenau Instituut (2020a). *Valued at work. Limits to digital monitoring at the workplace using data, algorithms and AI*. Den Haag: Rathenau Instituut

Rathenau Instituut (2020b). *Meer grip op digitalisering. Een internationale vergelijking van parlementaire werkvormen*. Den Haag: Rathenau Instituut

Rathenau Instituut (2019). "Zo brengen we AI in de praktijk vanuit publieke waarden." <https://www.rathenau.nl/nl/digitale-samenleving/zo-brengen-we-ai-de-praktijk-vanuit-europese-waarden>

Rathenau Instituut (2017a) *Urgent Upgrade. Protect public values in our digitized society*. Den Haag: Rathenau Instituut.

Rathenau Instituut (2017b) *Human rights in the robot age. Challenges arising from the use of robotics, artificial intelligence, and virtual and augmented reality*. Den Haag: Rathenau Instituut.

Schermer, B, J. van Ham en K. W. Falkena (2020). *Onvoorziene effecten van zelflerende algoritmen*. Amsterdam: Considerati.

TNO (2019) *Quick scan AI in de publieke dienstverlening* (eindpresentatie 8 april 2019). <https://www.rijksoverheid.nl/documenten/rapporten/2019/04/08/quick-scan-in-de-publieke-dienstverlening>.

Tweede Kamer der Staten-Generaal (2021a). *Klem tussen balie en beleid*. Rapport van de Tijdelijke Commissie Uitvoeringsorganisaties (TCU).

Tweede Kamer der Staten-Generaal (2021b). *Ongekend Onrecht*. Rapport van de Parlementaire ondervragingscommissie Kinderopvangtoeslag (POK).

Tweede Kamer der Staten-Generaal (2020). *Update vereist. Naar meer parlementaire grip op digitalisering*. Rapport van de Tijdelijke Commissie Digitale Toekomst (TCDT).

Van Eck, M., M. Bovens en S. Zouridis (2018). "Algoritmische rechtstoepassing in de democratische rechtsstaat." *Nederlands juristenblad* 93(40), 3008-3017.

Vetzo, M.J., J. Gerards en R. Nehmelman (2018). *Algoritmes en grondrechten*. Den Haag: Boom Juridisch.

Widlak, A. & R. Peeters (2018). *De digitale kooi. Over (on)behoorlijk bestuur door informatiearchitectuur. Of: hoe we de burger weer centraal zetten in een digitaliserende overheid*. Den Haag: Boom Bestuurskunde.

Wieringa, M. (2020). "What to account for when accounting for algorithms. A systematic literature review on algorithmic accountability." In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*: 1-18.

Wolswinkel, J. (2020a). "AR meets AI. AI: Een bestuursrechtelijk perspectief op een nieuwe generatie besluitvorming." *Computerrecht*, 2020(1): 22-29.

Wolswinkel, J. (2020b). *Willekeur of algoritme? Laveren tussen analoog en digitaal bestuursrecht*. Tilburg University.

WRR (2017). *Weten is nog geen doen – Een realistisch perspectief op redzaamheid*. Den Haag: WRR.

WRR (2016). *Big data in een vrije en veilige samenleving*. Den Haag: WRR.

Zouridis, S. (2000). *Digitale disciplineren. Over ICT, organisatie, wetgeving en het automatiseren van beschikkingen* (diss. Tilburg). Delft: Eburon 2000.

Zuiderveen Borgesius, F. (2018). *Discrimination, artificial intelligence, and algorithmic decision-making*. Strasbourg: Council of Europe, Directorate General of Democracy.