

Artificial Intelligence and Public Service Delivery

The use of AI in Allocation of Social Benefits

Policy Brief No. 5 /2020

Author
Moses Otsieno, Researcher

1. Introduction

Artificial Intelligence (AI)-systems are increasingly entering the realm of public services all over the world. One of the most famous examples of it is the Chinese 'social credit system', which uses wide-ranging types of personal data to allow or deny access to public services.¹ However, more and more countries across the globe are turning to AI-systems to assist in the allocation, management or termination of social benefits for citizens. These types of AI-systems perform or assist in the governmental task of adjudication: the granting of benefits or rights to citizens.² To illustrate, the Government of Netherlands has long used an algorithm to detect child care benefits fraud.³ The machine learning algorithm filtered out 'high fraud risk' applications that were then subjected to human monitoring. For the monitoring personnel, the system was a black-box. It was trained with examples of 'correct' and 'fraudulent' applications, using multiple data indicators, including information on migration backgrounds. The system – falsely - labelled a disproportionate amount of people with a migration background as fraudsters and on this basis the Authority– wrongfully – reclaimed thousands of euros of received benefits.⁴ Other countries also using AI-systems to assess eligibility for social security benefits include Sweden⁵ and Spain⁶. While there are advantages to the use of AI-systems for this purpose, this algorithmic turn in social benefits allocation also presents many ethical, legal and societal challenges. It is of utmost importance to adequately identify these challenges in order to mitigate risks.

2. Benefits of using AI for social benefits allocation

The use of AI-systems for social benefits distribution has several advantages. If we think just of the technical aspects, AI may be able to accelerate the decisions by implementing the rules at a large scale and with little "effort". This means public servants spend less time on the task, saving resources or allocating them to higher-importance tasks. From the point of view of the applicants, speed is also important, as they're not in a position to wait without harm. AI may also broaden the scope of the criteria for allocation, as it can take decisions based on more data points. If trained for that, it could also lead to fraud prevention and detection. Thus, one could argue that the use of AI would lead to fairer allocations, with less corruption or biases. Last, some would argue that AI can serve a predictive function, which could possibly improve fraud surveillance.

¹ Qian Sun, 'China's social credit system was due by 2020 but is far from ready', algorithmwatch.org

² Misuraca, G., and van Noordt, C., Overview of the use and impact of AI in public services in the EU, EUR 30255 EN, Publications Office of the European Union, Luxembourg, 2020.

³ [Björn ten Seldam & Alex Brenninkmeijer](#), 'The Dutch benefits scandal: a cautionary tale for algorithmic enforcement', eulawenforcementlaw.com 2020.

⁴ Autoriteit Persoonsgegevens, Belastingdienst/Toeslagen: De verwerking van de nationaliteit van aanvragers van kinderopvangtoeslag 2020.

⁵ Ibid

⁶ Nicolas Kayser-Bril, 'Spain: Legal fight over an algorithm's code', algorithmwatch.org 2019.

3. Societal challenges

However, in spite of the advantages, various ethical, legal and technical challenges arise. Foremost, AI-systems in the sphere of social benefits distribution limit several fundamental rights, mainly in relation to rights to privacy, non-discrimination and effective remedy. These fundamental rights limitations can be connected to problems arising in the fields of ethics, politics and socio-economics. Fundamental rights are not absolute: the law generally allows for limitation of fundamental rights for protection of economic wellbeing. Thus, AI-systems in the context of social benefits allocations may be allowed under certain conditions, even when this limits fundamental rights. The use of AI in public services in fact requires a balancing exercise between the public interest and the individual rights and interests.

3.1 Privacy

The first challenge concerns *the right to private life*, including the *right to data protection*, as protected under national frameworks on privacy and data protection in Africa and many other countries. Privacy does not only include the right to be left alone, but is also a precondition for the enjoyment of other rights, such as the right to self-determination and freedom of speech. Individuals under constant surveillance lack the freedom to develop personhood, and can be manipulated into making certain choices. AI-systems such as the Netherlands child benefits-system impact on the right to private life because the combining of different data sources allows the government to create a detailed (but not necessarily accurate) profile about one's personal life, most often without their knowledge, therefore also affecting ethical notions of individual freedom and autonomy (e.g. capacity to make free choices). These systems also limit the right to data protection, as AI-systems depend on large amounts of – often personal – data.

From a technical perspective, it is challenging that AI-systems evolve around enormous amounts of personal data. While large amounts of data may strengthen values such as equality and solidarity, when resulting in more accurate allocation of scarce resources. However, the focus should be on the quality, not quantity, of data. From a political (philosophy) perspective, it may be problematic that social benefits AI-systems encourage 'data hunger', creating an incentive to collect more private data, conflicting with privacy values. This relates to the classic problem of balancing public interests with individual ones, but in case of privacy, violations may also lead to a 'crisis of public trust', which eventually impacts state governance. Whether limitation of privacy rights is proportional in the case of social benefits allocation AI-systems depends on a whole range of factors, such as the type of data being processed, the legal basis, security measures, and combination with other data sets or sources. In the case of the Dutch child care benefits system, some of the personal data collected (namely migration backgrounds) were not relevant for the purpose of the system, and therefore disproportional.

3.2 Discrimination

The second challenge is related to *the right to non-discrimination*, as protected under Article 2 of the African Charter on Human and Peoples Rights and many African Constitutions. The problem with AI-systems is that they often produce biased results, mirroring or even amplifying the existing bias, most likely against certain protected categories, such as a specific race or gender. In the case of AI-powered social benefits allocation systems, it leads to the denial of benefits to specific groups, exacerbating historical inequalities. This can be caused by the data they learn on, the design of the algorithm or by

the way in which the system is used. From a technical perspective, the first two are difficult to solve, whilst the latter should be the object of broader policy and political considerations.

Discrimination is conflicting with ethical notions of human dignity (e.g. intrinsic value of all human beings) and fairness (e.g. equitable distribution of good and harm). AI-systems pose a risk of discrimination, targeting specific population groups more than others. Relatedly, on the socio-economic front, when faulty decisions are suggested by the system, the impact on lower-income households will be profound, while such families may have higher barriers to contest the decision. Use of the system when targeting individuals of lower socio-economic status would require intense scrutiny, oversight and verification (requiring high explainability standards) and provide low-barrier possibilities for recourse. Social benefits AI-systems with substantial negative effects for specific groups are disproportionate when balanced against the interest of administrative efficiency.

3.3 Effective remedy and explainability

The right to effective remedy may also be affected. First, every citizen affected by a government decision has the right to challenge it (administratively / in court). The framework needs to provide for the situation where the decision generated or influenced by the AI-system is not necessary final. Second, if the AI system is of a black-box nature, it may be difficult to assess the validity of the decision process. The burden of proof of the legality of the decision lies with the public entity or administration, which should be able to demonstrate reasonable grounds for the decision, based on the statutory criteria, as provided by the fundamental right to good administration. Public service institutions are under the legal and moral obligation to disclose the reasons underlying their decisions, which in this case means documentation on the programming and training methods, processes and techniques used to build, test and validate the AI system. This explains the way in which the algorithmic process contributed to the decision and locates potential bias that could lead to discrimination. The role of the human civil servant, which may range from a thorough and critical evaluation of the recommendation to a mere rubber-stamping, should also be clear, especially in the case of AI-assisted social benefits allocation systems.

3.4 Political considerations

The use of such a system for the specified use case entails a political choice. It greatly increases the possibilities of the state to control individuals, in a way that could not be foreseen when existing legislation was drafted. It's also politically significant to choose to assess benefits attribution instead of, for instance, target tax fraud in richer segments of the population. AI in social security and taxation may lead to shifts in the balance between state power and individual freedom, with the chances of having greater impact on those with less means to contest the decisions. Proportionality is therefore crucial: the aim of ensuring that public resources are justly allocated should be balanced against the infringements on the personal lives of citizens and a violation of democratic norms such as equal dignity and an equal stance in the collective decisions taken.

4. An expert system for social security benefits allocation?

It is unlikely that the advantages of AI-systems for social security benefits allocation outweigh the societal issues and fundamental rights limitations. Possibly, 'classical AI' expert systems would be more suitable. Unlike a contemporary machine learning system based on neural networks and fed with big data, an expert system would rely on a well-defined model and clearly specified procedural rules,

mitigating several of the risks discussed above. In general, expert systems would allow for direct explainability of the reasoning process, making it easier for a “human in the loop” to spot mistakes, and improve possibilities for recourse. In the “safest” case, the model could be based on the existing procedural legislation and policy for social benefits. As an input, the system would only use personal data directly and lawfully related to the criteria for obtaining social benefits, such as the officially declared income. Even if additional, circumstantial data were used, the system would retain a high degree of explainability, allowing for better scrutiny and easier redress. This would substantially improve issues associated with privacy, discrimination and effective remedy.

5. Conclusion

The use of AI-Systems for the allocation of social benefits can be challenged from several standpoints. We believe there are circumstantial and intrinsic reasons for rejecting its use. Currently, due to technical problems related to bias and explainability, it would be hard to build such a system without significant flaws and without infringing fundamental rights. However, even if such technical problems would be overcome, its political justification would face problems related to legitimacy and proportionality. Some more ‘humanity’ in social benefits allocation, e.g. in the form of expert systems, would make social security truly ‘social’.