

The use of AI can disproportionately negatively affect marginalised groups in our society. Is there any benefit to using AI in our legal system?

By Amelia Twining

Is Artificial Intelligence (“AI”) “the last invention that humanity will ever need to invent,” as claimed by Swedish philosopher Nick Bostrom¹, or does the reality lie closer to tech entrepreneur Elon Musk’s warning that with “artificial intelligence we’re summoning the demon?”²

Certainly, when viewed from the perspective of marginalised groups in our society today, such as women, racial minorities or people from lower socio-economic groups, Musk’s view resonates. This is because AI has the potential to not only perpetuate, but also amplify existing bias and discrimination. There are several reasons for this. Firstly, the algorithms that underpin AI are often based on data sets that are themselves not diverse or representative. For example, Amazon programmed an AI tool to save them time and money by sifting job applicants. Only later did they realise their system was biased against women applying for software development roles, because the system had concluded that the higher number of male applicants – reflecting the industry’s existing gender bias – was evidence that female applicants were less likely to be successful in the role, thereby reinforcing an existing societal bias. Unsurprisingly, the tool was dropped in 2017³.

A further reason that AI can lead to discriminatory outcomes is because of conscious and unconscious bias by those involved in coding AI. A notorious example is facial recognition software, which according to a 2018 study by MIT Media Lab⁴, was 99% accurate when asked to recognise a photo of a white man, but increasingly inaccurate the darker the skin, with an error rate of up to 35% for darker skinned women. This arguably reflects the fact that the tech industry’s workforce is mostly white, with only 2.5% of Google’s employees being black according to New York University’s AI Now Institute⁵, thereby embedding racial prejudice into both the data sets selected for the AI and the way the system is trained.

¹ <https://nickbostrom.com/ethics/ai>

² <https://www.washingtonpost.com/news/innovations/wp/2014/10/24/elon-musk-with-artificial-intelligence-we-are-summoning-the-demon/>

³ <https://www.bbc.co.uk/news/technology-45809919>

⁴ <https://www.media.mit.edu/articles/facial-recognition-software-is-biased-towards-white-men-researcher-finds/>

⁵ <https://ainowinstitute.org/discriminatingystems.pdf>

AI can also have a negative impact on marginalised groups, by taking away their employment. This is because as the sophistication of AI grows, so does its ability to perform low-skilled tasks that were previously done by humans. McKinsey & Co estimate that between 400 and 800 million jobs will be displaced in this way by 2030⁶. This trend disproportionately impacts marginalised groups who are more likely to be performing these sorts of roles today, leading to higher rates of job losses and poverty in an already disadvantaged group.

However, while AI can disproportionately negatively affect marginalised groups, it also has the potential to revolutionise the legal system for the benefit of all, by improving its efficiency, accuracy, fairness and objectivity as well as broadening access to justice.

For example, AI's ability to process and analyse large amounts of data much faster and more accurately than humans opens up the opportunity of e-discovery, whereby gigabytes of data can be rapidly searched for keywords, duplicate documents or connections that would otherwise involve hours of billable work to complete. An example would be due diligence for a corporate transaction, where AI can help lawyers by analysing documents and then highlighting references to areas such as litigation issues or key contract clauses for further human scrutiny. In a recent experiment, twenty experienced lawyers were pitted against an AI programme and tasked with reviewing five non-disclosure agreements and identifying within these thirty legal issues such as arbitration and indemnification. While it took the human lawyers 92 minutes on average to complete the task, the AI took only 26 seconds⁷. This automation of time-consuming and often repetitive tasks, frees up lawyers to focus on higher value activity, thereby reducing costs, reducing delays and speeding up the legal process.

A further benefit of using AI in the legal system is its greater level of accuracy. In the study cited above, for example, the AI achieved 95% accuracy, versus the lawyers who on average only achieved an 85% accuracy rate. Further, AI has the ability to spot patterns and relationships in often vast data sets that might otherwise have been missed by human review. This can be particularly useful in accurately identifying relevant case law and statutes that might otherwise have been missed. To this end, the Incorporated Council of Law Reporting for England and Wales has launched an AI tool called CaseGenie, that enables users to upload a document or text and then "identify otherwise seamlessly unconnected or novel cases."⁸ The benefit of this is that it helps lawyers find what Donald Rumsfeld famously referred to as the "unknown unknowns" – the previously unknown precedent cases or judgements that might help strengthen a case and deliver a better outcome for clients.

Another advantage of using AI in the legal system is the opportunity it provides for fairness. For example, in a recent study of 100,000 judicial decisions in the US, researchers found that some judges released more than 90% of defendants on bail, while others released only 50%. AI can in theory eliminate these biases and ensure consistency in decision making, ensuring that all defendants are treated fairly, rather than have decisions influenced by their colour, religion or sexual orientation. A practical illustration

⁶ <https://www.mckinsey.com/featured-insights/future-of-work/jobs-lost-jobs-gained-what-the-future-of-work-will-mean-for-jobs-skills-and-wages>

⁷ <https://www.superlegal.ai/blog/aiyslwyer/>

⁸ <https://www.artificiallawyer.com/2021/11/02/iclr-launches-ai-driven-case-genie/>

of this is the use of AI to make sentencing recommendations based on a defendant's criminal history and other relevant factors. A case in point is the trial of Eric Loomis in Wisconsin, USA, who was jailed for six years in 2013, having pleaded guilty to the charge of eluding a police officer. In determining his sentence, the judge considered both his criminal record and a score, provided by an AI system called Compas, that predicted that Loomis posed a high risk of re-offending⁹. While Loomis subsequently unsuccessfully appealed the decision to the US Supreme Court, on the basis that the use of a tool whose inner workings could not be examined violated due process, institutions such as the Brookings Institute continue to argue for AI to play a greater role in sentencing¹⁰. One recent simulation by the US National Bureau of Economic Research concluded that AI sentencing programs could be used to cut crime across by up to 25% or reduce jail populations by up to 42 % without any increase in crime rates¹¹.

A final potential benefit of using AI in the legal system is that it can improve access to justice. A November 2020 survey by the UK's Legal Services Board revealed that more than three million consumers per year encounter a need for legal advice that went unmet – unsurprising, given that 87% of those surveyed believe they cannot afford it¹². AI can, for example, provide automated legal guidance, making legal services more accessible and affordable to a wider range of people, including those who cannot afford a lawyer or do not have access to legal services. Additionally, AI can be used to create online legal services that are available 24/7, making legal help more accessible to people who live in rural areas or who cannot take time off work to visit a lawyer.

Of course, while there are potential benefits of the growing use of AI in the legal system, the risks to marginalised groups remain. For example, in the UK AI has been used for “predictive policing,” by analysing historic data to predict where and when certain crimes might happen¹³. In the US, AI is used to help determine whether a defendant should be released on bail. The problem is that these tools are not neutral and suffer from the same biases in data sets and patterns of discrimination previously outlined. The human rights organisation Liberty argues that this encourages “an approach to policing based on discriminatory profiling¹⁴”, while the MIT Technology Review maintains that AI is “sending people to jail – and getting it wrong¹⁵.”

Whatever we may or may not think about the risks and benefits of AI, one thing is clear - it is here to stay. Indeed, according to Bill Gates, AI is “as fundamental as the creation of the microprocessor, the personal computer, the Internet, and the mobile phone¹⁶”. This is not to deny the clear evidence that AI can disproportionately impact marginalised groups in our society or ignore the risk that this might be allowed to perpetuate. But the significant potential benefits to the use of AI in the legal system – not to mention its

⁹ https://en.wikipedia.org/wiki/Loomis_v._Wisconsin

¹⁰ <https://www.brookings.edu/blog/techtank/2017/07/20/its-time-for-our-justice-system-to-embrace-artificial-intelligence/>

¹¹ <https://www.cs.cornell.edu/home/kleinber/w23180.pdf>

¹² https://legalservicesboard.org.uk/wp-content/uploads/2020/11/The-State-of-Legal-Services-Narrative-Volume_Final.pdf

¹³ <https://www.lawsociety.org.uk/topics/research/algorithm-use-in-the-criminal-justice-system-report>

¹⁴ <https://www.libertyhumanrights.org.uk/fundamental/predictive-policing/>

¹⁵ <https://www.technologyreview.com/2019/01/21/137783/algorithms-criminal-justice-ai/>

¹⁶ <https://www.bbc.co.uk/news/technology-65032848>

inexorable rise across every aspect of our lives - suggests that we need to find a way to implement and regulate the use of AI in a way that is transparent, inclusive and fair. On a practical level, this means ensuring that AI systems are developed, trained, and tested in an ethical and fair manner by conducting fairness audits, implementing algorithmic accountability measures, and ensuring diversity and inclusivity in the development teams. Rather than simply reject the use of AI out of hand, therefore, the focus should be on harnessing its potential for the benefit of all, without exception.