

CIPR

CHARTERED INSTITUTE
OF PUBLIC RELATIONS

AI and data ethics

Skills Guide

CIPR #AlinPR Primer series

PROFESSIONAL
DEVELOPMENT
cipr.co.uk

CPD 

This guide is worth 5 CPD points

Contents

1	Learning outcomes	3
2	How AI uses data	3
3	Data collection	4
4	Big data	4
	Ethical implications	
5	Bias	4
6	Privacy, control and transparency	5
7	Glossary	6
8	References and further reading	6
9	Biography	7
10	Acknowledgements	7

AI and data ethics

Ben Verinder, FCIPR, Dip CIPR, Founding Chartered PR Practitioner, MMRS

There are a wide range of topics that touch on the matter of ethics relating to artificial intelligence (AI), including economic welfare, equality, national security and AI 'rights'.

This primer focuses on AI's use of data. It is intended to explain:

- How AI uses data
- The ethical implications of that use.

Learning outcomes

In this guide you will learn:

- What constitutes artificial intelligence
- The types of data AI uses
- How AI use of data raises ethical questions for organisations.

Artificial intelligence enables software to perform difficult tasks by learning through repetition, rather than following a sequence of instructions. AI is used to solve complex problems, often involving a large amount of data, efficiently, inexpensively and at scale.

How AI uses data

AI commonly involves machine learning, in which automated methods (algorithms) find patterns in data to generate insight and make the decisions and predictions that power AI.

Data used by AI systems is typically separated into different categories:

- Training data is a set of examples that teaches the AI to do its job. Through a series of cycles, the AI analyses the data to assess how patterns within the variables (the details within it) relate to the task at hand
- Validation data, new and unfamiliar to the AI, is used to select and fine tune the model
- Testing data is used to evaluate how the final model, after all the improvements, will perform in the real world.

The data will be rich in detail, capable of improving the AI after many training cycles. The more variables within the data, the faster the algorithms can learn and fine-tune their predictive analyses.

The more data analysed, the better the AI gets at spotting anomalies, probabilities and patterns – speech patterns that make it easier to sound like a human, for instance, behaviour patterns to identify fraud or visual patterns to detect objects on a street.

Success, therefore, depends on comprehensive, large data sets. It also depends on the relevance of information. If data is not clean (i.e. up-to-date, accurate) and properly

organised, the AI model won't be able to establish the right patterns and may not be able to perform the task at hand.

Data collection

Because AI typically needs such a lot of data, companies developing AI often supplement their own data with that purchased from external sources (data from different sources is called triangulated data). External sources can include subscription suppliers, public and private crowdsourcing of data, as well as new services which generate synthetic data for training purposes.

There is also a growing demand for services that help companies label more complex data such as images and video.

It can be difficult at the outset to estimate the amount of training data that will be needed as it depends on the type of machine learning employed, the number and characteristics of the features selected, the quality of the owned data and the degree of accuracy a model needs for a task to be performed.

Big data

Big data is a type of data set with properties that lend themselves to training AI, as #AlinPR panel's Vice Chair Kerry Sheehan explains:

"The term 'big data' is used frequently, and quite often incorrectly. Big data is described as having three contributing characteristics: it must have high velocity, high variety, and high volume. Distilled down to its essence, big data accumulates quickly, in a variety of types, and there is a lot of it. Processing it surpasses the capabilities of most traditional software.

"Most PR professionals and communicators aren't routinely working with what could truly be defined as big data yet. This is beginning to change as communicators become more aware of how corporate/organisational data are being used to inform everything from sales strategies and logistical planning to service take up. The volume of news and information being created globally does qualify as big data, but the amount of that news that is relevant to a specific company or institution is a subset of that – so, in most cases it is not technically big data even if there is high volume."

Ethical implications

Bias

Incomplete or inaccurate training data can lead to bias, influencing high-stakes decisions and potentially reinforcing existing social biases. There is evidence of flaws and biases in facial recognition software powered by AI, for instance. In April 2019 a group of artificial intelligence researchers – including those at Google, Facebook and Microsoft – called on Amazon to stop selling its facial-recognition technology to law enforcement agencies because, they argue, it is demonstrably biased against women and people from black and minority ethnic groups.

Bias in this area is commonly the result of incomplete or imbalanced data sets. Because most systems operate by comparing a person's face to a range of images in

a large database, if the database isn't broad enough and contains mostly or only male Caucasian faces, for instance, that's what the programme will learn to recognise.

The way in which data is collected and handled can also lead to bias in the results. Human biases can affect the way data is classified or what is excluded from the data set. If data is crowd-sourced, scraped, or otherwise gathered from digital applications it may well be biased by virtue of the types of people who use online services. These people – typically the socio-economically advantaged – will commonly be over-represented in the data.

AI can also be deliberately used to amplify bias. Insurance companies purposefully discriminate in order to minimise perceived risk, as do banks when they make loans; AI systems are now being used to support underwriting decisions and mortgage applications.

Bias embedded into algorithms is hard to spot because we cannot necessarily see into the 'black box' of decision making. For example, an autonomous vehicle developed by American technology company Nvidia went on the roads in 2016, without anyone knowing how it made its driving decisions. Given this difficulty of understanding exactly how AI systems have reached a given decision, bias is often only revealed when the AI is deployed. This was the case with a system used to make sentencing decisions in courts across the United States which over-estimated recidivism rates among African-American and Hispanic men.

Privacy, control and transparency

There is a necessary tension between AI's voracious appetite for data, which can help reduce bias, and the principles of minimisation of personal data and transparency in data processing.

Public awareness of data collection, storage and sharing issues has risen and expectations of how data is handled have changed as a result of legislation such as the General Data Protection Regulation (GDPR), and in light of infamous data breaches such as that involving Facebook and Cambridge Analytica.

Regulations and user expectations are not limited to privacy. They include:

- Control over how data is used
- Limits on the duration and nature of data storage
- Better transparency, in terms of what data is being collected and with whom it is being shared.

There are those who argue that AI and GDPR are fundamentally incompatible. At the heart of GDPR (and equivalent data protection regulations in other regions, such as the California Consumer Privacy Act) is the principle of informed consent. But the nature of machine learning systems means it is typically not possible to see how their decisions have been reached. How can we consent to something that we can't understand? And how can we withdraw consent once our data has been deployed?

AI also poses a challenge to our desire to limit the range of data collected about us and the length of time it is retained. The amount of data needed for AI to progress, this argument goes, is going to increase and this, in turn, boosts the chances of people's data being collected, stored and manipulated without their consent or knowledge.

Glossary

Artificial intelligence: The ability of software to perform difficult tasks by learning through repetition, rather than following a sequence of instructions. AI is used to solve complex problems, often involving a large amount of data, efficiently, inexpensively and at scale.

Big data: data sets characterised by high velocity, variety and volume.

Bias: inclination or prejudice for or against one person or group, especially in a way considered to be unfair.

GDPR: General Data Protection Regulation, approved by the EU Parliament on 14 April 2016 and enforced on 25 May 2018, intended to protect EU citizens' data privacy, harmonize data privacy laws across Europe, and regulate how organisations across the region approach data privacy.

References and further reading

Elgan, M. Artificial intelligence needs your data, all of it, ComputerWorld, 22 February 2016 www.computerworld.com/article/3035595/artificial-intelligence-needs-your-data-all-of-it.html

Ismail, N. The success of artificial intelligence depends on data, Information Age, 23 April 2018 www.information-age.com/success-artificial-intelligence-data-123471607/

Ismail, N. Is business data AI compatible?, Information Age, 19 June 2017 www.information-age.com/business-data-ai-compatible-123466827/

Lorica, B. Data collection and data markets in the age of privacy and machine learning, O'Reilly, 18 July 2018 www.oreilly.com/ideas/data-collection-and-data-markets-in-the-age-of-privacy-and-machine-learning

The Norwegian Data Protection Authority, Artificial intelligence and privacy, January 2018 www.datatilsynet.no/globalassets/global/english/ai-and-privacy.pdf

The Open Data Institute, Whitepaper: Helping organisations navigate ethical concerns in their data practices, 13 September 2017 www.scribd.com/document/358778144/ODI-Ethical-Data-Handling-2017-09-13

Smith, D. What is AI Training Data?, Gengo A, 29 January 2019 <https://gengo.ai/articles/what-is-ai-training-data/>

Data collection

Hardinges, J. What is a Data Trust, The Open Data Institute <https://theodi.org/article/what-is-a-data-trust/>

Korolov, M. AI's biggest risk factor: Data gone wrong, COI US www.cio.com/article/3254693/ais-biggest-risk-factor-data-gone-wrong.html

Melendez, C, Data is the lifeblood of AI, but how do you collect it?, Infoworld, 28 August 2018 www.infoworld.com/article/3296044/data-is-the-lifeblood-of-ai-but-how-do-you-collect-it.html

Bias

AI Now Institute, AI Now 2017 Report, New York University, November 2017
https://ainowinstitute.org/AI_Now_2017_Report.pdf

Beall, A. It's time to address artificial intelligence's ethical problems, Wired, 24 August 2018 www.wired.co.uk/article/artificial-intelligence-ethical-framework

Brackeen, B. Facial Recognition Software is Not Ready for Use by Law Enforcement, Tech Crunch, June 25, 2018 <https://techcrunch.com/2018/06/25/facial-recognition-software-is-not-ready-for-use-by-law-enforcement>

Chapman, L and Brustein, J. A.I. Has a Race Problem, Bloomberg Businessweek, July 2, 2018. www.bloomberg.com/news/articles/2018-06-26/ai-has-a-race-problem

Metz, C and Singer, N. A.I. Experts Question Amazon's Facial-Recognition Technology, New York Times, April 3 2019
www.nytimes.com/2019/04/03/technology/amazon-facial-recognition-technology.html

Nehme, A. The Short-Term Ethical Concerns of AI, Towards Data Science, June 6 2018
<https://towardsdatascience.com/the-short-term-ethical-concerns-of-ai-c201c03bc0ac>

Pegoraro, R. Microsoft Argues Facial-Recognition Tech Could Violate Your Rights, Yahoo, July 27, 2018 <https://finance.yahoo.com/news/microsoft-argues-facial-recognition-tech-violate-rights-182552188.html>

Privacy

Beall, A. It's time to address artificial intelligence's ethical problems, Wired, 24 August 2018 www.wired.co.uk/article/artificial-intelligence-ethical-framework

Information Commissioners Office, Big data, artificial intelligence, machine learning and data protection, 4 September 2017 <https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf>

Matthias Avocats, Artificial Intelligence and the GDPR: how do they interact?, 27 November 2017
www.avocats-mathias.com/technologies-avancees/artificial-intelligence-gdpr

Meyer, D. AI Has a Big Privacy Problem and Europe's New Data Protection Law Is About to Expose It, Fortune, May 25 2018
<http://fortune.com/2018/05/25/ai-machine-learning-privacy-gdpr/>

Richard, L, AI and the Future of Business, Sunderland Software City, 5 April 2019
<http://insights.sunderlandsoftwarecity.com/post/102fh5n/ai-and-the-future-of-business>

Biography

Ben Verinder is managing director of Chalkstream, a reputation research, insight and management consultancy. He is a CIPR Founding Chartered Practitioner and lead chartered status assessor who began his PR career in technology and software.

THANKS

With many thanks to Kerry Sheehan, Emma Thwaites, members of the #AlinPR panel and CIPR Fellows for their suggestions and contributions.

Last updated: July 2019