



Machine learning: the power and promise of computers that learn by example

SUMMARY

Machine learning is a branch of artificial intelligence that allows computer systems to learn directly from examples, data, and experience.

Recent years have seen exciting advances in machine learning. Systems which only a few years ago performed at noticeably below-human levels can now outperform humans at some specific tasks, and many people now interact with systems based on machine learning every day, for example in image recognition systems on social media; voice recognition systems, used by virtual personal assistants; and recommender systems, such as those used by online retailers.

As the field develops further, machine learning shows promise of supporting potentially transformative advances in a range of areas, and the social and economic opportunities which follow are significant. In healthcare, machine learning is creating systems that can help doctors give more accurate or effective diagnoses for certain conditions. In transport, it is supporting the development of autonomous vehicles, and helping to make existing transport networks more efficient. For public services it has the potential to target support more effectively to those in need, or to tailor services to users. And in science, machine learning is helping to make sense of the vast amount of data available to researchers today, offering new insights into biology, physics, the social sciences, and more.

The Royal Society's report *Machine learning: the power and promise of computers that learn by example* sets out the actions necessary to allow us to benefit fully from the development of machine learning and to address some of the associated challenges.

To read the full report, visit royalsociety.org/machine-learning

Extracting value from data

The recent success of machine learning owes no small part to the explosion of data that is available in some areas, such as image or speech recognition. This data has provided a vast number of examples, which machine learning systems can use to improve their performance.

Supporting the development of machine learning requires an amenable data environment, based on:

- continued Government open data efforts;
- new models of data sharing that respect privacy and enable carefully-managed access to certain datasets, for example from the NHS;
- resources for data management within research funding; and
- extending the lifecycle of open data through open data standards.

It has been estimate that

90%

of the world's data has been created within the last five years.

The direct value of public sector data has been estimated at

£1.8 billion

with the wider social and economic benefits from this totalling

£6.8 billion

23%

of the UK population lack basic digital skills.

Creating value from machine learning

SKILLS

As machine learning systems become ubiquitous, and a more significant part of people's lives and livelihoods, three skills needs follow.

Firstly, as daily interactions with machine learning become the norm, a basic understanding of the use of data and these systems will become an important tool required by people of all ages and backgrounds. Introducing key concepts in machine learning as well as some of the key social and ethical issues at school can help cultivate these skills.

Secondly, to ensure that a wide range of sectors and professions have the absorptive capacity to use machine learning in ways that are useful for them, new mechanisms are needed to create a pool of informed users or practitioners. This requires adjusting university course provision in disciplines such as law, healthcare, and finance, to introduce machine learning to future professionals across these fields. A new funded programme of Masters courses may also help to increase the number of informed users of machine learning.

Thirdly, further support is needed to build advanced skills in machine learning. There is already high demand for people with advanced skills, and additional resources to increase this talent pool are critically needed. These resources include increasing provision for training PhD students, and creating mechanisms to recruit and retain outstanding research leaders in machine learning in the academic sector.

BUSINESS

There is a vast range of potential benefits from further uptake of machine learning across industry sectors, and the economic effects of this technology could play a central role in helping to address the UK's productivity gap.

Businesses of all sizes across sectors need to have access to appropriate support that helps them to understand the value of data and machine learning to their operations. Such support includes:

- access to talent;
- advice via government support mechanisms for business; and
- measures to promote machine learning through the industrial strategy.

Machine learning and society

While offering potential for new businesses or areas of the UK economy to thrive, the disruptive potential of machine learning brings with it challenges for society, and questions about its social consequences.

Some of these challenges relate to the way in which new uses of data reframe traditional concepts of, for example, privacy or consent. While it is not appropriate to set up governance structures for machine learning per se, governance surrounding the use of data requires a new framework to keep pace with the challenges in the 21st century.

Other challenges stem from the relationship between humans and machine learning systems. The Royal Society carried out the UK's first in-depth public dialogues on machine learning, and found that attitudes towards machine learning, both positive and negative, vary depending on the circumstances in which it is being used.

Continuous engagement between machine learning researchers and the public will be important as the field develops. This should be complemented by relevant ethics training for machine learning postgraduates.

A new wave of research

Machine learning is a vibrant field of research, with a range of exciting areas for further development across different methods and applications. There is a collection of specific research questions where progress would directly address potential public concerns around machine learning, or constraints on its wider use:

- **Interpretability:** Can we create powerful machine learning systems where the reasons for particular decisions or recommendations can be understood or interrogated?
- **Verification and validation:** Can we create more advanced, and accurate, methods of verifying machine learning systems?
- **Privacy:** What are the technical solutions that can maintain the privacy of datasets, while allowing them to be used in new ways?
- **Fairness and dealing with real-world data:** How can real-world data be curated into usable forms, addressing ‘real-world’ messiness, and systemic – or social – bias?
- **Causality:** How can machine learning methods discover cause-effect relationships?
- **Human-machine interaction:** How do we design machine learning systems so they can work with humans safely and effectively?
- **Security:** How do we ensure machine learning systems are not vulnerable to cyber-attack?

Support for research in these areas is needed to help ensure continued public confidence in the deployment of machine learning systems.

Opening further debates

In addition to these near-term policy measures, machine learning raises further questions for the future:

- What impact will machine learning have on work, and how should we manage this?
- How can we ensure the field develops in a direction that provides broad social benefits?
- How can we make sure the benefits of machine learning are shared?

Ensuring the best possible environment for the safe and rapid deployment of machine learning will be essential for enhancing the UK's economic growth, wellbeing, and security, and for unlocking the value of 'big data'.

Careful stewardship will be necessary to help ensure that the benefits from machine learning are shared across society.

Action is needed in key areas – shaping the data landscape, building skills, supporting business, maintaining public confidence, and advancing research.

“Machine learning will have an increasing impact on our lives and lifestyles over the next five to ten years. There is much work to be done so that we take advantage of machine learning’s potential and ensure that the benefits are shared, especially as this could be a key area of opportunity for the UK in the coming years.”

Professor Peter Donnelly FMedSci FRS,
Chair of the Royal Society Working Group
on Machine Learning, and Director of the
Wellcome Trust Centre for Human Genetics,
University of Oxford.

The Royal Society

The Royal Society is a self-governing Fellowship of many of the world's most distinguished scientists drawn from all areas of science, engineering, and medicine. The Society's fundamental purpose, as it has been since its foundation in 1660, is to recognise, promote, and support excellence in science and to encourage the development and use of science for the benefit of humanity.

The Society's strategic priorities are:

- Promoting excellence in science
- Supporting international collaboration
- Demonstrating the importance of science to everyone

For further information

The Royal Society
6 – 9 Carlton House Terrace
London SW1Y 5AG

T +44 20 7451 2500

W royalsociety.org



Founded in 1660, the Royal Society is the independent scientific academy of the UK, dedicated to promoting excellence in science.

Registered Charity No 207043
Issued: April 2017 DES4702_2

Cover image

© shulz.