Big Data and A.I.
-Ethical and societal implications-

Prof. Roberto V. Zicari

Frankfurt Big Data Lab

www.bigdata.uni-frankfurt.de

UCBerkeley, October 16, 2018

Data as an Economic Asset

ﷺ

"I think we're just beginning to grapple with implications of data as an economic asset"

–Steve Lohr *(The New York Times)*

ﷻ Data has become a new economic asset.

ﷻ The companies with big data pools can have great economic power.

What is more important, vast data pools, sophisticated algorithms or deep pockets?

"No one can replicate your data. It's the defensible barrier, not algorithms."

-- **Andrew Ng**, Stanford professor

Algorithms and Data

ॐ

"AI is akin to building a rocket ship. You need a huge engine and a lot of fuel. The rocket engine is the learning algorithms but the fuel is the huge amounts of data we can feed to these algorithms."

-- Andrew Ng

It is important to note that Big Data is of NO use unless it is analyzed.

Interplay and implications of Big Data and Artificial Intelligence

"In marketing and advertising, a decision that is better on average is plenty good enough. You've increased sales and made more money. You don't really have to know why."
--Steve Lohr (*New York Times*).

ℛ The Big Data revolution has made the recent AI advances possible.

ℛ The initial use of Big Data technologies started with Marketing. One stage in the life cycle of an emerging science, marketing is a low-risk – and, yes, lucrative.

AI and Big Data
for higher-stakes decisions

ௐ Technology is moving beyond increasing the odds of making a sale, to being used in higher-stakes decisions like medical diagnosis, loan approvals, hiring and crime prevention…

Example:
Personalised Healthcare
Digital Biomarkers

෴

*" …our application of mobile and sensor technology to monitor symptoms, disease progression and treatment response – the so called "**Digital Biomarkers**". We have our most advanced programmes in Multiple Sclerosis (MS) and Parkinson`s Disease (PD), with several more in development. Using these tools, a longitudinal real-world profile is built that, in these complex syndromes, helps us to identify signals and changes in symptoms or general living factors, which may have several potential benefits."*

— Bryn Roberts
Global Head of Operations for **Roche Pharmaceutical Research & Early Development**

෴   **On using AI and Data Analytics in Pharmaceutical Research. Interview with Bryn Roberts** ODBMS Industry Watch, September 10, 2018

Example: Healthcare
Democratisation of data

*" I'm intrigued by the general trend towards empowering individuals to share their data in a secure and controlled environment. **Democratisation of data** in this way has to be the future. Imagine what we will be able to do in decades to come, when individuals have access to their complete healthcare records in electronic form, paired with high quality data from genomics, epigenetics, microbiome, imaging, activity and lifestyle profiles, etc., supported by a platform that enables individuals to share all or parts of their data with partners of their choice, for purposes they care about, in return for services they value – very exciting! "*

 *– **Bryn Roberts***

Example:
Transportation: *Railigent*

ઝ Transportation is using a lot of AI-driven algorithms

ઝ *Railigent* is a cloud based platform, by Siemens, designed to help rail operators and rail asset owners, to improve fleet availability and improve operations, for example by **enabling intelligent data gathering, monitoring, and analysis for prescriptive maintenance in the rail transport industry.**

ઝ Railigent contains a cloud based platform layer to support ingest and storage of large and diverse data sets, high end data analytics and applications. This layer is open, both for customers and partners.

ઝ On top of this layer, Railigent provides a large set of applications for monitoring and analyzing rail assets. Also here applications and components can be provided by partners or customers. Target is to help customers improve fleet availability and improve operations.

Automate or **Augment humans**?

ℰ

ℰ You can use AI technologies either to **automate** or to **augment** humans.

In the first case, machines replace people, in the second case machine complements people (at least in theory).

Implications

ೞ When we use Big Data in domains such as healthcare, decisions are **practically** and **ethically** very different then marketing and advertising.

ೞ These are crucial decisions about **individual people's lives.** Better on average isn't good enough.

ೞ For these kinds of decisions, issues of **accuracy, fairness and discrimination** come into play.

New ethical and legal questions

Some AI applications may raise new **ethical** and **legal** questions, for example related to liability or potentially biased decision-making.

For example, algorithms are used to review loan applications, recruit new employees and assess potential customers, and **if the data are skewed** the decisions recommended by **such algorithms may be discriminatory against certain categories or groups.**

Example : Healthcare Predictive Insights

"What happens if my algorithm is wrong? Someone sees the wrong ad. What's the harm? It's not a false positive for breast cancer."

*-- **Claudia Perlich**,* Data Scientist

*"Citizens and businesses alike need to be able to **trust** the technology they interact with, and have effective safeguards protecting fundamental rights and freedoms.*

*In order to increase **transparency** and minimise the risk of bias, AI systems should be developed and deployed in a manner that allows humans to understand the basis of their actions.*

Explainable AI *is an essential factor in the process of strengthening people's trust in such systems."*

-- Roberto Viola

Director General of DG CONNECT (Directorate General of Communication Networks, Content and Technology) at the European Commission.

"I think the most important aspect of this question is the simple **acknowledgement** that **intelligent system designers do have ethical responsibilities.**"

--John Markoff (The New York Times)

"We have a **profound ethical responsibility to design systems** that have a positive impact on society, obey the law, and adhere to our highest **ethical** standards."

–Oren Etzioni
Chief Executive Officer **Allen Institute for Artificial Intelligence.**

"We are moving towards changes in how ethics has to be perceived: away from individual decisions with specific and knowable outcomes, towards actions by many unaware that they may have taken actions with unintended consequences for anyone."

*--**Andrej Zwitter** (Dean Faculty Campus Fryslân / Prof. of Governance and Innovation)*

"I think ethical software development for AI is not fundamentally different from ethical software development in general. The interesting new question is: **when AIs learn by themselves, how do we keep them from gowing astray?**

Fixed rules of ethics, like Asimov's three laws of robotics, are too rigid and fail easily. (That's what his robot stories were about.) **But if we just let machines learn ethics by observing and emulating us, they will learn to do lots of unethical things.**

So maybe AI will force us to confront what we really mean by ethics before we can decide how we want AIs to be ethical."

--Pedro Domingos *(Professor at University of Washington)*

5 core principles to keep AI ethical

(UK Select Committee on Artificial Intelligence)

    AI must be a force for good - and diversity

    Intelligibility and fairness

    Data protection

    Flourishing alongside AI

    Confronting the power to destroy

Source: World Economic Forum
https://www.weforum.org/agenda/2018/04/keep-calm-and-make-ai-ethical/

How to Understand Decisions?

ﻌﻫ　At present we do not really understand how Advanced AI-techniques such as used in Deep learning (e.g. neural networks) really works.

ﻌﻫ　This is due to the technical complexity of such advanced neural networks, which need huge amount of data to learn properly. It is a try and error.

ﻌﻫ　This poses an ethical and societal problem:

What if the decision made using AI-driven algorithm harmed somebody, and you cannot explain how the decision was made?

"Since the algorithms learn from data, it's not as easy to understand what they do as it would be if they were programmed by us, like traditional algorithms. But that's the essence of machine learning: that it can go beyond our knowledge to discover new things. A phenomenon may be more complex than a human can understand, but not more complex than a computer can understand.

*And in many cases we also don't know what humans do: for example, we know how to drive a car, but we don't know how to program a car to drive itself. **But with machine learning the car can learn to drive by watching video of humans drive."***

--- Pedro Domingos

"Debugging an AI is more like domesticating an animal than debugging a program."

*" Debugging AI systems is harder than debugging traditional ones, but not impossible. Mainly it requires a different mindset, that allows **for non determinism and a partial understanding** of what's going on**. Is the problem in the data, the system, or in how the system is being applied to the data?***

-- Pedro Domingos

Autonomous systems

Is it essential to **keep a human in the loop**?

If not, do you run the **risk of abdicating ethical responsibility for system design?**

Is it realistic? If something can be partially automated, will it eventually be fully automated?

Example : Autonomous Cars

ଔ

Let`s consider an autonomous car that relies entirely on an algorithm that had taught itself to drive by watching a human do it.

What if one day the car crashed into a tree, or even worse killed a pedestrian?

Autonomous Cars (cont.)

෯ This is referred in the literature as the *crash assignment,* especially between automated vehicles and non automated vehicles.

෯ Some researchers have indicated that automated vehicles will need to be programmed with some sort of ethical system in order to make decisions on how to crash. We will come back to this point later when we talk about *Ethics by Design*.

Autonomous Cars (cont.)

���� Few studies however, have been conducted on how particular ethical theories will actually make crash decisions and how these ethical paradigms will affect automated vehicle programming. [Wesley 2015]

Autonomous Cars (cont.)

&

"If the learning took place before the car was delivered to the customer, the car's manufacturer would be liable, just as with any other machinery. The more interesting problem is if the car learned from its driver.

Did the driver set a bad example, or did the car not learn properly?"

--**Pedro Domingos**

Autonomous Cars (cont.)

In the area of autonomous vehicle, there are initiatives to define Ethical principles, spanning from independent associations such as:

The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems , which aims [IEEE 2018] :

 "To ensure every stakeholder involved in the design and development of autonomous and intelligent systems is educated, trained, and empowered to prioritize ethical considerations so that these technologies are advanced for the benefit of humanity"

(source https://standards.ieee.org/develop/indconn/ec/autonomous_systems.html)

Autonomous Cars (cont.)

to political bodies such as the German *Ethics Commission on Automated Driving.*

"*The Ethics Commission on Automated and Connected Driving has developed initial* **guidelines for policymakers and lawmakers** *that will make it possible to approve automated driving systems but that set out special requirements in terms of safety, human dignity, personal freedom of choice and data autonomy.*"
-- Prof. Di Fabio

(source
https://www.bmvi.de/SharedDocs/EN/PressRelease/2017/084-ethic-commission-report-automated-driving.html)
The Ethics Commission's complete report can be found here:
www.bmvi.de/report-ethicscommission

The Ethics Commission's report comprises 20 propositions.

The key elements are: (from the web site : https://www.bmvi.de/SharedDocs/EN/PressRelease/2017/084-ethic-commission-report-automated-driving.html)

- Automated and connected driving is an ethical imperative if the systems cause fewer accidents than human drivers (positive balance of risk).
- Damage to property must take precedence over personal injury. In hazardous situations, the protection of human life must always have top priority.
- In the event of unavoidable accident situations, any distinction between individuals based on personal features (age, gender, physical or mental constitution) is impermissible.
- In every driving situation, it must be clearly regulated and apparent who is responsible for the driving task: the human or the computer.
- It must be documented and stored who is driving (to resolve possible issues of liability, among other things).
- Drivers must always be able to decide themselves whether their vehicle data are to be forwarded and used (data sovereignty).

Closing the Gap
Designers of intelligent systems

ള Are computer system designers (i.e. Software Developers, Software Engineers, Data Scientists, Data Engineers, etc,), the ones who will decide what the impact of these technologies are and whether to replace or augment humans in society?

Accountable AI?

罗

ଔ **Do we need some sort of auditing tool**?
The technology has to be able to explain itself, to explain how a data-driven algorithm came to the decision or recommendation that it did. Is it technically feasible?

ଔ How much **Transparency** is desired/ possible?

ଔ Do we wish **"Human in the loop"** for most of these kinds of decisions for the foreseeable future?

How is it possible to define incentives for using an ethical approach to software development, especially in the area of AI?

Ethics by Design

&

"Solving actual moral problems is not simply a matter of choosing the "best" of several possible responses. It is also a matter of devising possible responses. Design practice in engineering affords important lessons about addressing practical problems" [Whitbeck 1996]

Ethics by Design

CS

*"Ethics by Design concerns the methods, algorithms and tools needed to endow **autonomous agents with the capability to reason about the ethical aspects of their decisions**, and the methods, tools and formalisms to guarantee that an agent's behavior remains within given moral bounds. In this context some questions arise: How and to what extent can agents understand the social reality in which they operate, and the other intelligences (AI, animals and humans) with which they coexist? What are the ethical concerns in the emerging new forms of society, and how do we ensure the human dimension is upheld in interactions and decisions by autonomous agents?. **But overall, the central question is: "Can we, and should we, build ethically-aware agents?"***

Source: Ethics by Design: necessity or curse?
http://www.aies-conference.com/wp-content/papers/main/AIES_2018_paper_68.pdf

ࢯ Data, AI and Intelligent systems are becoming sophisticated tools in the hands of a variety of stakeholders, including political leaders.

Human Motivations

ℭℬ

"It is absolutely essential that we control the machines, and every indication is that we will be able to do so in the foreseeable future. I do worry about human motivations too. Someone said: I'm not worried about robots deciding to kill people, I'm worried about politicians deciding robots should kill people."
-- **Oren Etzioni**

Good Intention vs. Good Data

ﻭ *Good data* reflects reality and thus can help us gain insights into how the world works. That does not make such discovery ethical, even though the discover is correct.

ﻭ *Good intentions* point towards an ethical use of data, which helps protect us again unethical data uses, but does not prevent false big data analysis.

This is a long way of saying we need both, albeit for different reasons.

--**Viktor Mayer-Schönberger**, *Professor of Internet Governance and Regulation at Oxford University (UK).*

The thing that **motivates** my **actions** will determine the **direction** I am going

The Individual and Collective Conscience

The **individual and collective conscience** is the existential place where the most significant things happen.

Research, Change, Decision and Choice can take two diametrically opposite directions:

can be either "pro or contra" the human person".

Source: "The good society and the future of jobs: Can solidarity and fraternity be part of business decisions?" MAY 8 -10, 2014 – Vatican City

"Data for Humanity"

ঝ An initiative I started with **Andrej Zwitter** (Professor at the University of Groningen) at the end of 2015, with the goal to:

*bring people and institutions together who share the **motivation to use data for the common good.***

which calls for the use of **five ethical principles for the use of data**- has reached more than 1,000 signatories worldwide.

www.bigdata.uni-frankfurt.de/dataforhumanity/

Five ethical principles when using data

෩ Do no harm

෩ Use data to help create peaceful coexistence

෩ Use data to help vulnerable people and people in need

෩ Use data to preserve and improve natural environment

෩ Use data to help create a fair world without discrimination

Do we need to regulate the development of artificial intelligence?

"While self-regulation can be a first stage in applying an ethical approach, public authorities must ensure that the regulatory framework that applies to AI technologies is fit for purpose and in line with our values and fundamental rights.

For example, the Commission is currently assessing the safety and national and EU liability frameworks in light of the new challenges, and we will examine whether any legislative changes are required."

-- Roberto Viola
Director General of DG CONNECT (Directorate General of Communication Networks, Content and Technology) at the European Commission.

Regulatory Frameworks and AI

ℭℨ

"The Commission closely monitors all relevant developments related to AI and, if necessary, we will review our existing legal framework. The EU has a strong and balanced regulatory framework to build on in order to develop a sustainable approach to AI technologies. This includes high standards in terms of safety and product liability, EU-wide rules on network and information systems security and stronger protection of personal data that came into force in May 2018."

 --**Roberto Viola**

Regulations vs. Innovation

"I think the biggest challenge is that in the rail business we have a very large set of old and country specific regulations that date back many decades. **These regulations are meant to protect passengers, but some of them are not anymore fitting to the modern capabilities of technology and instead drive cost and slow innovation down dramatically."**

-- Gerhard Kress (*Director Mobility Data Services Siemens*)

AI vision for the future

*"It should be all of us. Right now it is mainly the companies that have lots of data and sophisticated machine learning systems, but all of us **– as citizens and professionals and in our personal lives – should become aware of what AI is and what we can do with it.***

How can I use AI to do my job better, to find the things I need, to build a better society?

*Just like driving a car does not require knowing how the engine works, but it does require knowing how to use the steering wheel and pedals, **everyone needs to know how to control an AI system, and to have AIs that work for them and not for others**, just like they have cars and TVs that work for them."*

-- Pedro Domingos