The BSI PAGIT Project:
The role of standards in the governance of innovative technologies

Standards Matter Conference

Sheraton Hotel, Edinburgh, 22 June, 2017

Joyce Tait
Director, Innogen Institute, University of Edinburgh

innogen

The Open University

THE UNIVERSITY of EDINBURGH

The process of regulatory ‘capture’

- Cell therapies treated as ‘drugs’ for regulatory purposes
- GM crops treated as ‘a plant pest’ (US regulations)
- GM fish treated as ‘a drug’ (US Regulations)
- GM organisms regulated according to the process by which they were produced, not the properties of the product (EU regulatory system)

These are all disruptive innovations

Disruptive and Incremental Innovation

Incremental innovation:

- Enables stepwise improvements in a company’s current innovation system, creating competitive advantage within the same sector without challenging the prevailing business models.
- There is usually a clear regulatory precedent

Disruptive innovation:

- Cannot be accommodated within a company’s current business model. It needs new areas of R&D; new modes of production; new routes to market, and sometimes new markets and industry sectors.
- There is either no clear regulatory precedent or no agreement on the choice of regulatory system

innogen

Tait, J. (2007) Systemic Interactions in Life Science Innovation. *Technology Analysis and Strategic Management, 19(3),* 257-277, May 2007*.*

Disruption depends on the sector

An innovation that is path-breaking for one industry sector can be path dependent for another.

- **GM crops**
- **Cell therapies**

Advice to policy makers and regulators:
Thinking about how to regulate a new highly innovative technology – decide for which industry sector that innovation would be most path-dependent and least disruptive of existing business models and choose the regulatory system under which they currently operate.

Tait, J. (2007) Systemic Interactions in Life Science Innovation. *Technology Analysis and Strategic Management, 19(3),* 257-277, May 2007*.*

Linking the concepts of disruption and innovation

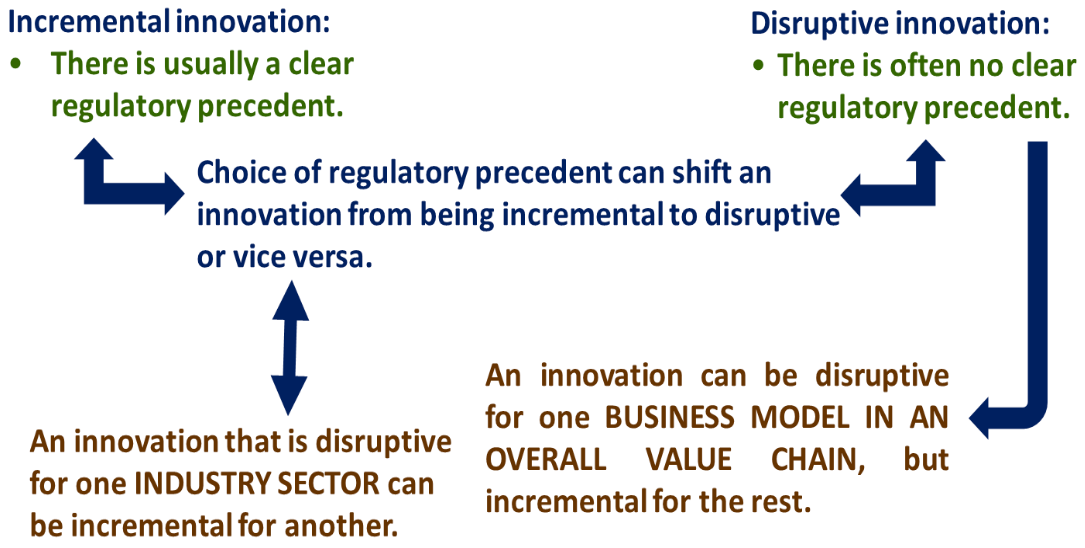

Incremental innovation:

- There is usually a clear regulatory precedent.

Disruptive innovation:

- There is often no clear regulatory precedent.

Choice of regulatory precedent can shift an innovation from being incremental to disruptive or vice versa.

An innovation that is disruptive for one INDUSTRY SECTOR can be incremental for another.

An innovation can be disruptive for one BUSINESS MODEL IN AN OVERALL VALUE CHAIN, but incremental for the rest.

Our regulatory systems need to change

- **The Innovation Principle**
  "… to improve the quality and application of EU legislation and as a result, to stimulate confidence, investment and innovation"

Plus

Two Regulatory Principles

- **Proportionality Principle**
- **Adaptation Principle**

innogen

PAGIT Framework

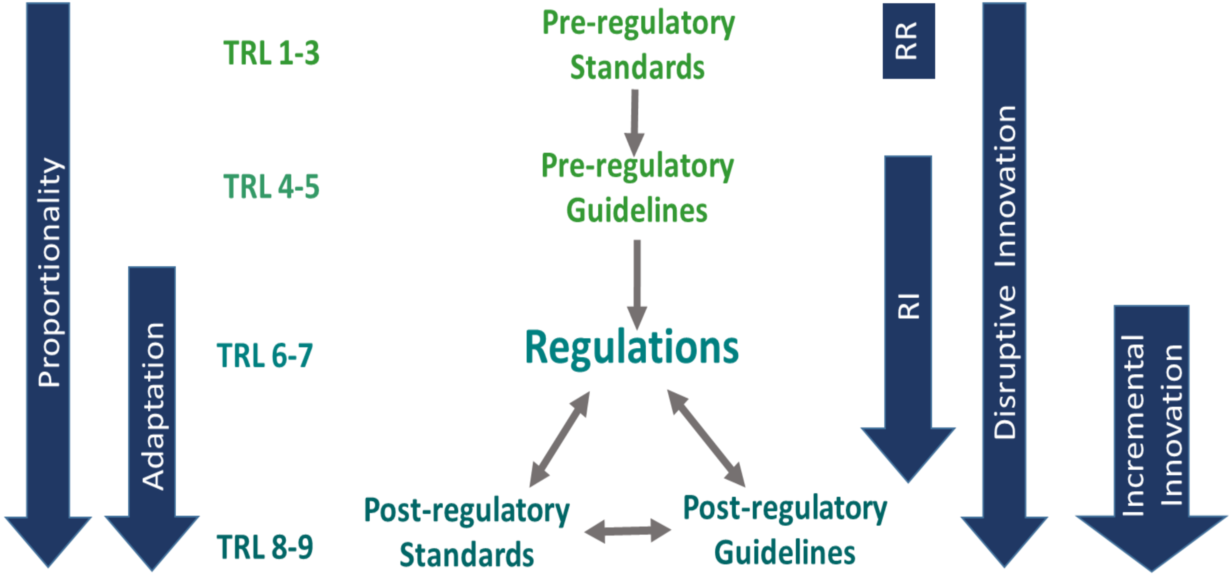

innogen

Responsible Research and Innovation

Approach includes compliance with the Innovation, Adaptation and Proportionality Principles

Proposes a more equitable allocation of responsibility

- All companies
  - Comply with Corporate Responsible Innovation Standard (based on ISO 26000)
- Where an innovation is disruptive or raises societal concerns:
  - Monitor the benefits and risks of the innovation throughout development
  - For stakeholder engagement, **all** stakeholders should comply with a Responsible Engagement Standard

innogen

The Brexit context

- Where and how far can the UK depart from current EU regulatory systems?
- UK Government's desire to lead internationally in developing a regulatory test-bed for innovative technologies
- Protecting our access to international markets – how far can regulatory adaptation go?
- The concept of 'regulatory equivalence'

innogen

Full Report

Proportionate and adaptive governance of innovative technologies

The role of regulations, guidelines and standards

Joyce Tait and Geoffrey Banda
Innogen Institute, University of Edinburgh

bsi.

Department for Business Innovation & Skills

innogen

http://www.bsigroup.com/research-pagit-uk

innogen

Thank you

innogen

GOVERNING THE AUTOMATED GOVERNORS

What does the GDPR say about automated decision-making and how can we regulate it?

BSI Standards Matter

Tristan Henderson

School of Computer Science
University of St Andrews
tnhh.org
tnhh@st-andrews.ac.uk

University of St Andrews

FOUNDED 1413

BEFORE WE BEGIN

- Disclaimer: IANAL
- Computer science academic at the University of St Andrews
- Law student at the University of Edinburgh

DECISIONS, DECISIONS EVERYWHERE

Predicting crime?

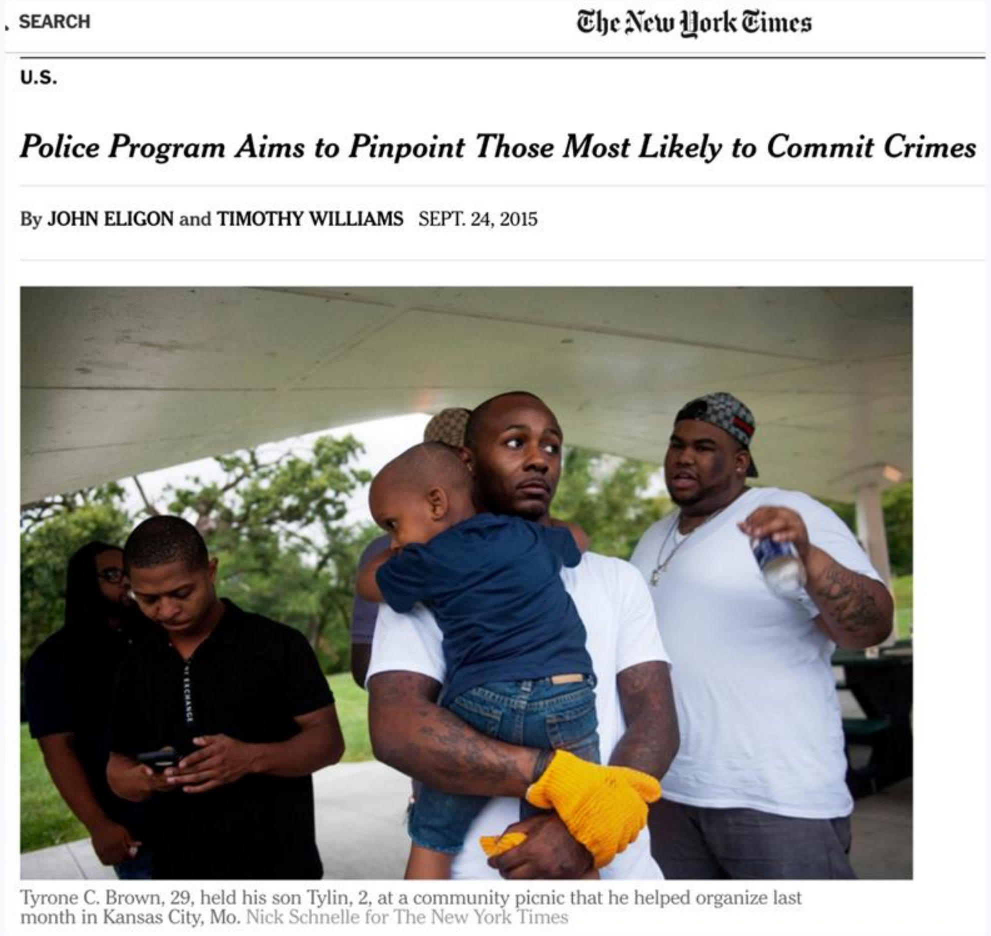

SEARCH

The New York Times

U.S.

Police Program Aims to Pinpoint Those Most Likely to Commit Crimes

By JOHN ELIGON and TIMOTHY WILLIAMS SEPT. 24, 2015

Tyrone C. Brown, 29, held his son Tylin, 2, at a community picnic that he helped organize last month in Kansas City, Mo. Nick Schnelle for The New York Times

nytimes.com/2015/09/25/us/police-program-aims-to-pinpoint-those-most-likely-to-commit-crimes.html

DECISIONS, DECISIONS EVERYWHERE

Predicting life choices?

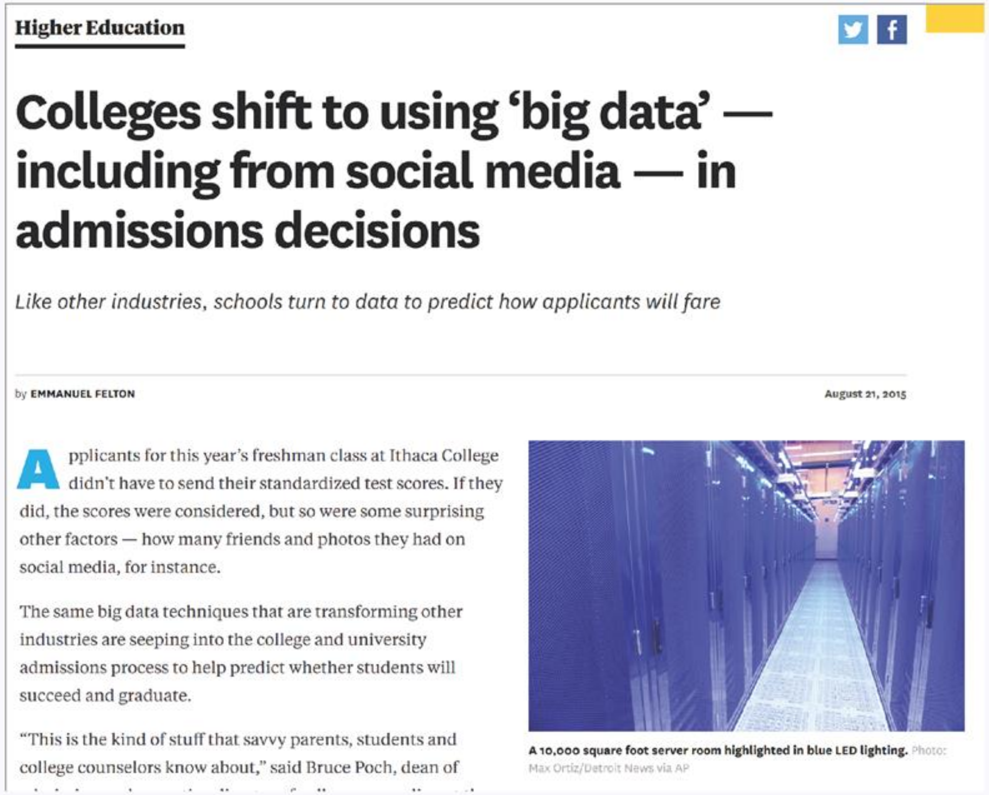

Higher Education

Colleges shift to using ‘big data’ — including from social media — in admissions decisions

Like other industries, schools turn to data to predict how applicants will fare

by EMMANUEL FELTON

August 21, 2015

Applicants for this year’s freshman class at Ithaca College didn’t have to send their standardized test scores. If they did, the scores were considered, but so were some surprising other factors — how many friends and photos they had on social media, for instance.

The same big data techniques that are transforming other industries are seeping into the college and university admissions process to help predict whether students will succeed and graduate.

“This is the kind of stuff that savvy parents, students and college counselors know about,” said Bruce Poch, dean of

A 10,000 square foot server room highlighted in blue LED lighting. Photo: Max Ortiz/Detroit News via AP

hechingerreport.org/colleges-shift-to-using-big-data-including-from-social-media-in-admissions-decisions/

DECISIONS, DECISIONS EVERYWHERE

Racial targeting?

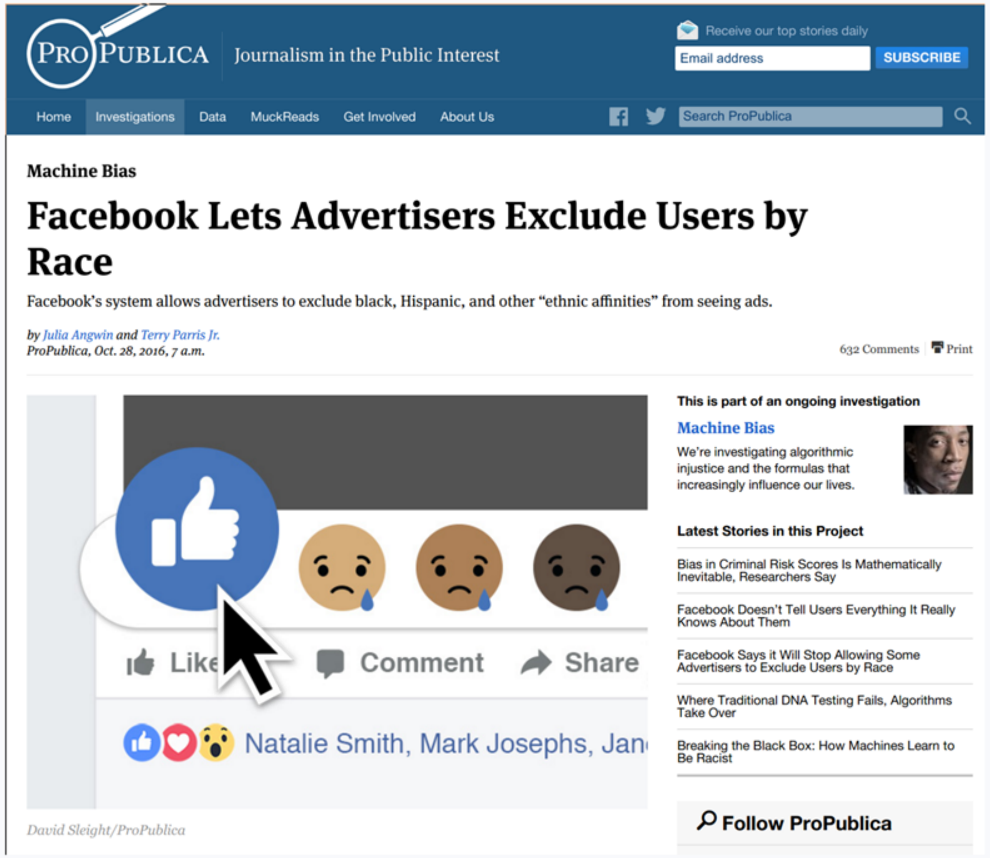

Machine Bias

Facebook Lets Advertisers Exclude Users by Race

Facebook's system allows advertisers to exclude black, Hispanic, and other "ethnic affinities" from seeing ads.

by Julia Angwin and Terry Parris Jr.
ProPublica, Oct. 28, 2016, 7 a.m.

632 Comments | Print

David Sleight/ProPublica

This is part of an ongoing investigation

Machine Bias

We're investigating algorithmic injustice and the formulas that increasingly influence our lives.

Latest Stories in this Project

Bias in Criminal Risk Scores Is Mathematically Inevitable, Researchers Say

Facebook Doesn't Tell Users Everything It Really Knows About Them

Facebook Says it Will Stop Allowing Some Advertisers to Exclude Users by Race

Where Traditional DNA Testing Fails, Algorithms Take Over

Breaking the Black Box: How Machines Learn to Be Racist

www.propublica.org/article/
facebook-lets-advertisers-exclude-users-by-race

DECISIONS, DECISIONS EVERYWHERE

Racist decisions?

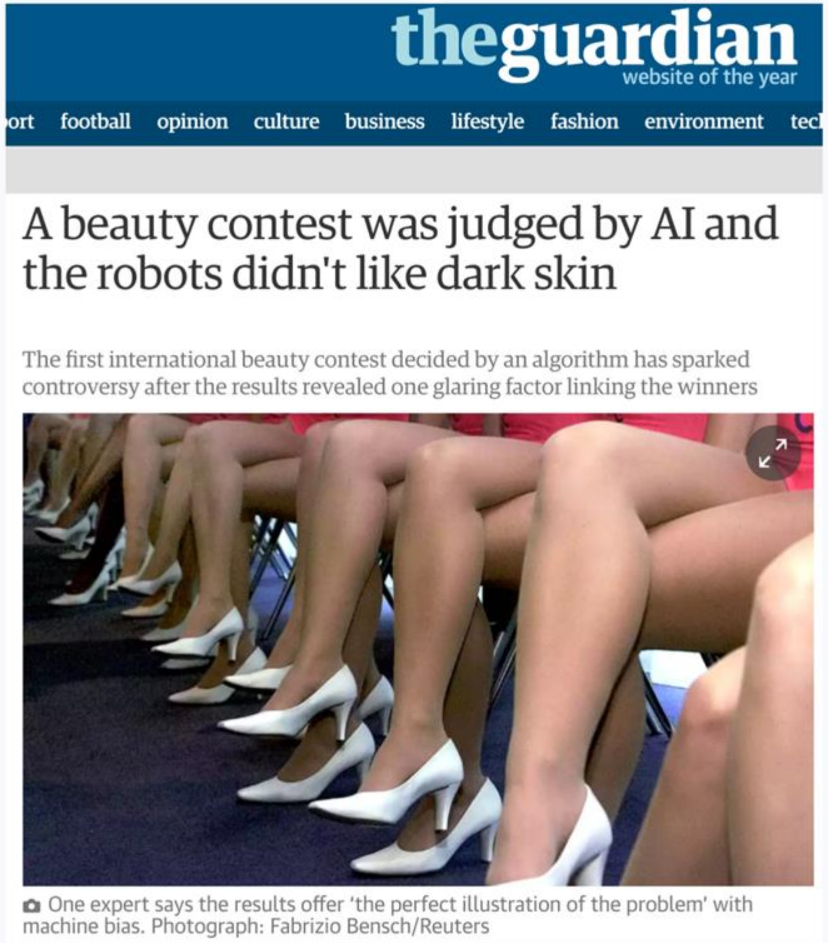
theguardian
website of the year
ort football opinion culture business lifestyle fashion environment tecl

A beauty contest was judged by AI and the robots didn't like dark skin

The first international beauty contest decided by an algorithm has sparked controversy after the results revealed one glaring factor linking the winners

One expert says the results offer 'the perfect illustration of the problem' with machine bias. Photograph: Fabrizio Bensch/Reuters

gu.com/p/527na

SHOULD WE BE CONCERNED?

- Automated decision-making is used everywhere
- The "data-driven society"[1] is increasingly enamoured with machine-learning classifiers
- Decisions about hiring, firing, crime, targeted news and adverts, beauty contests, ...
- These decisions can have serious societal, ethical and legal consequences
    - "weapons of math destruction"[2]
- What does the law have to say about it?
- Should the law do something about it?

[1] A. Pentland. The data-driven society. *Scientific American*, 309(4):78–83, 2013

[2] C. O'Neil. *Weapons of Math Destruction*. Allen Lane, 2016

GDPR ARTICLE 22

Automated individual decision-making, including profiling

1. The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.
2. Paragraph 1 shall not apply if the decision:
   (a) is necessary for entering into, or performance of, a contract between the data subject and a data controller;
   (b) is authorised by Union or Member State law to which the controller is subject and which also lays down suitable measures to safeguard the data subject's rights and freedoms and legitimate interests; or
   (c) is based on the data subject's explicit consent.
3. In the cases referred to in points (a) and (c) of paragraph 2, the data controller shall implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision.

GDPR ARTICLE 22

Automated individual decision-making, including profiling

1. The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.
2. Paragraph 1 shall not apply if the decision:
   (a) is necessary for entering into, or performance of, a contract between the data subject and a data controller;
   (b) is authorised by Union or Member State law to which the controller is subject and which also lays down suitable measures to safeguard the data subject's rights and freedoms and legitimate interests; or
   (c) is based on the data subject's explicit consent.
3. In the cases referred to in points (a) and (c) of paragraph 2, the data controller shall implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision.

GDPR ARTICLE 15

Right of access by the data subject

1. The data subject shall have the right to obtain from the controller confirmation as to whether or not personal data concerning him or her are being processed, and, where that is the case, access to the personal data and the following information:

   (h) the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject.

GDPR ARTICLE 15

Right of access by the data subject

1. The data subject shall have the right to obtain from the controller confirmation as to whether or not personal data concerning him or her are being processed, and, where that is the case, access to the personal data and the following information:
   (h) the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject.

DOES THE GDPR INTRODUCE A RIGHT TO EXPLANATION?

- Perhaps, say the non-lawyers[3]
- Perhaps not, say the lawyers[4]
- Art 22:
  - *solely* on automated processing (no human in the loop)
  - what are “legal effects” or “significantly affects”?
  - right is to obtain human intervention, not to explanation
  - which human can intervene?
- Art 15:
  - right to obtain confirmation
  - *ex post* or *ex ante* explanation? “envisaged consequences”
  - individual decision or overall logic of system?

[3] B. Goodman and S. Flaxman. European Union regulations on algorithmic decision-making and a “right to explanation”. In *2016 ICML Workshop on Human Interpretability in Machine Learning*, 2016

[4] S. Wachter, B. Mittelstadt, and L. Floridi. Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation, 2016

WHAT DOES THE TECHNOLOGY LOOK LIKE?

- **NOT** “algorithms”
- Just *classification* algorithms
- Types of machine learning classifiers used for prediction:
  - linear regression: easy to understand, low performance
  - Bayesian classifiers, decision trees: more difficult to understand, better performance
  - support vector machines, neural networks: very difficult to understand, even better performance
  - deep learning: even more difficult to understand, currently very popular
- All these techniques need *training data* to develop a predictive model:
  - supervised learning: training data are labelled ground-truth
  - unsupervised learning: algorithm develops clusters on its own

EXAMPLE DEEP LEARNING CLASSIFIER

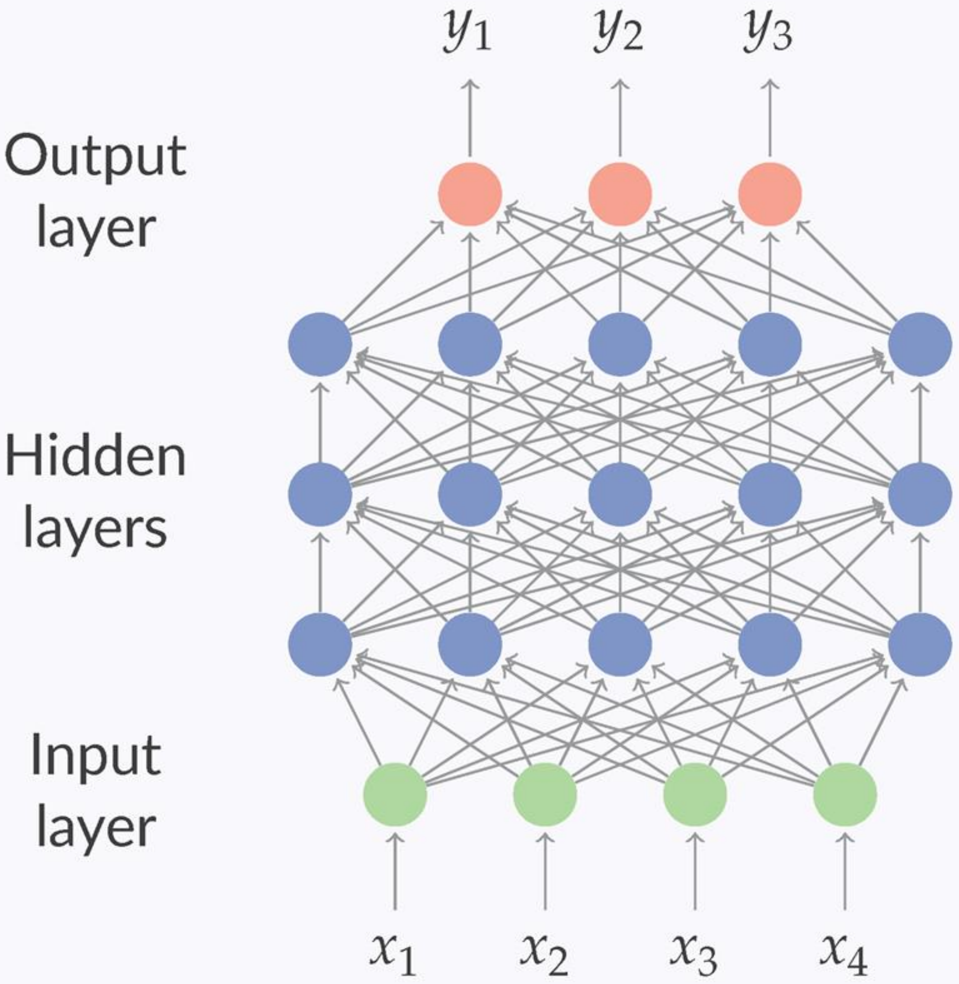

TECHNICAL SOLUTIONS

- Current relevant technical work can be broadly characterised as:
  - fair classifiers
  - interpretable classifiers
  - human-computer interaction / usability (not discussed here)
- Fairness / lack of discrimination for a given system; training set + test set + model
  - logic is determined by the training set
- Interpretation means auditing which features are useful, or changing parameters to monitor effect
  - not necessarily explaining a particular decision
- Does the GDPR's focus on protecting an *individual's* data prevent us from regulating a *collective* system?
  - a citizen may need access to an entire dataset, not just their own data, to achieve suitable redress

LEGAL SOLUTIONS

- How can we regulate decision-making using the law?
- Extend "meaningful information about the logic involved" (Arts 13-15)?
  - approach taken by Wachter *et al.*[5]
- Exercise the right to portability (Art 20)?
  - doesn't apply to inferred data according to A29WP[6]
- But there are other levers, where perhaps standards can help:
  - Data Protection by Design (Art 25)
  - Certification (Art 42)

[5] S. Wachter, B. Mittelstadt, and L. Floridi. Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation, 2016

[6] Article 29 Data Protection Working Party. Guidelines on the right to data portability, 2017-04-05. http://ec.europa.eu/newsroom/document.cfm?doc_id=44099

GDPR ARTICLE 42

Certification

1. The Member States, the supervisory authorities, the Board and the Commission shall encourage, in particular at Union level, the establishment of data protection certification mechanisms and of data protection seals and marks, for the purpose of demonstrating compliance with this Regulation of processing operations by controllers and processors. The specific needs of micro, small and medium-sized enterprises shall be taken into account.
3. The certification shall be voluntary and available via a process that is transparent.
6. The controller or processor which submits its processing to the certification mechanism shall provide the certification body referred to in Article 43, or where applicable, the competent supervisory authority, with all information and access to its processing activities which are necessary to conduct the certification procedure.

GDPR ARTICLE 42

Certification

1. The Member States, the supervisory authorities, the Board and the Commission shall encourage, in particular at Union level, the establishment of data protection certification mechanisms and of data protection seals and marks, for the purpose of demonstrating compliance with this Regulation of processing operations by controllers and processors. The specific needs of micro, small and medium-sized enterprises shall be taken into account.

3. The certification shall be voluntary and available via a process that is transparent.

6. The controller or processor which submits its processing to the certification mechanism shall provide the certification body referred to in Article 43, or where applicable, the competent supervisory authority, with all information and access to its processing activities which are necessary to conduct the certification procedure.

DEEP LEARNING AND THE GDPR

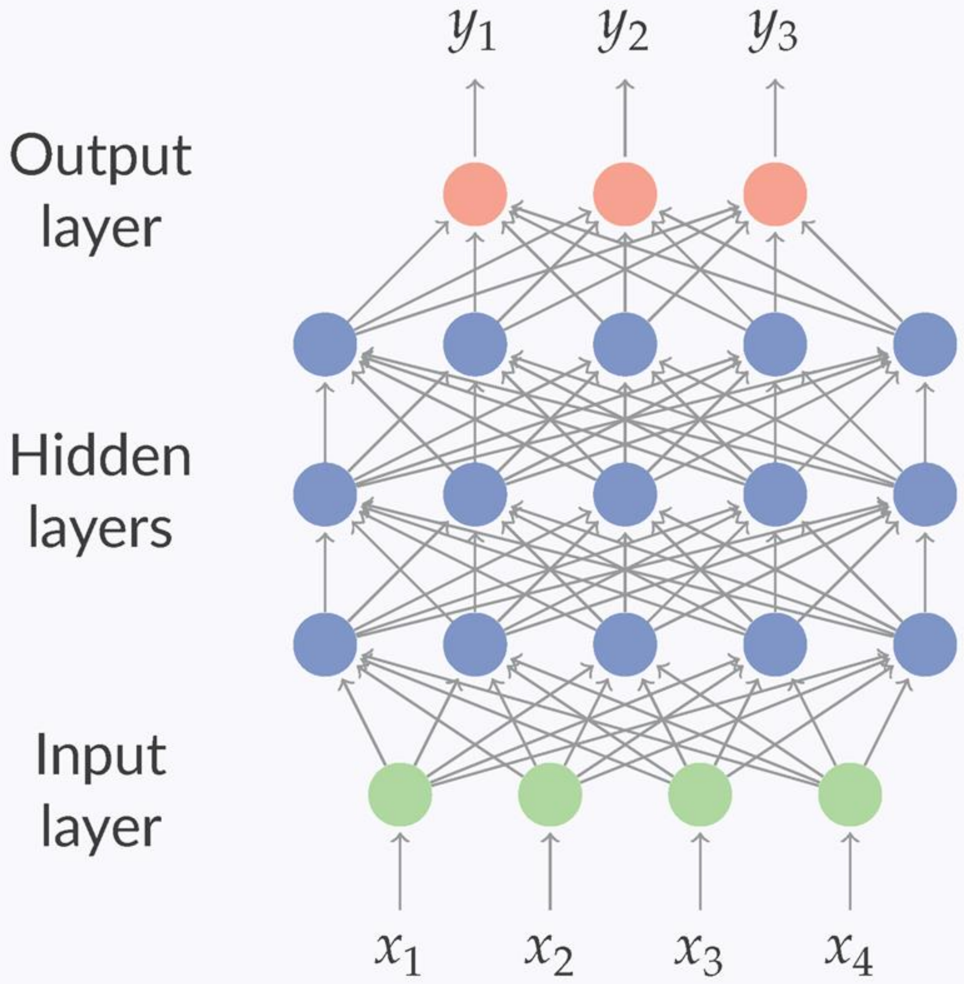

DEEP LEARNING AND THE GDPR

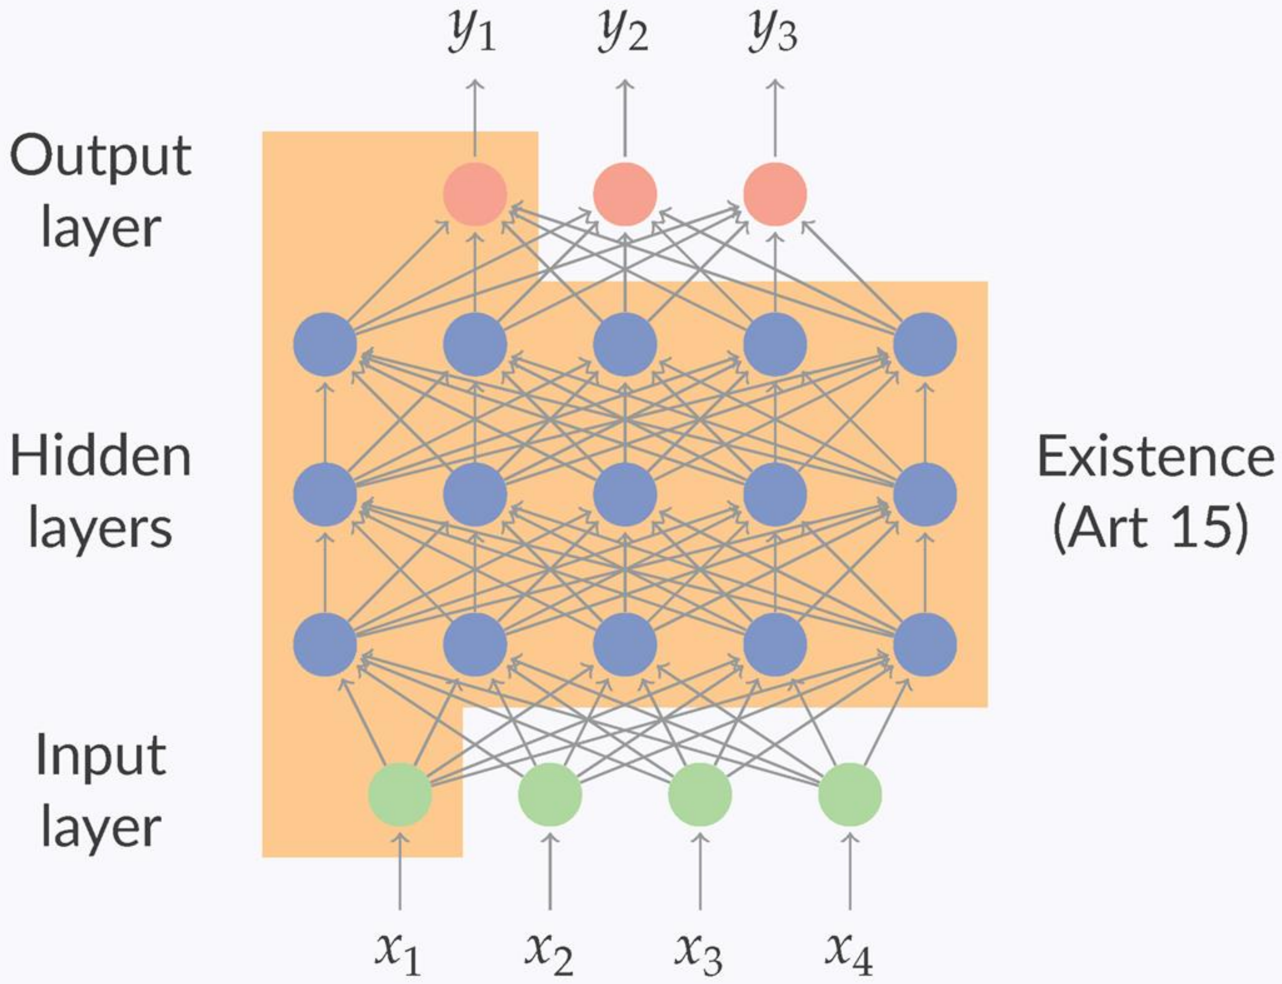

DEEP LEARNING AND THE GDPR

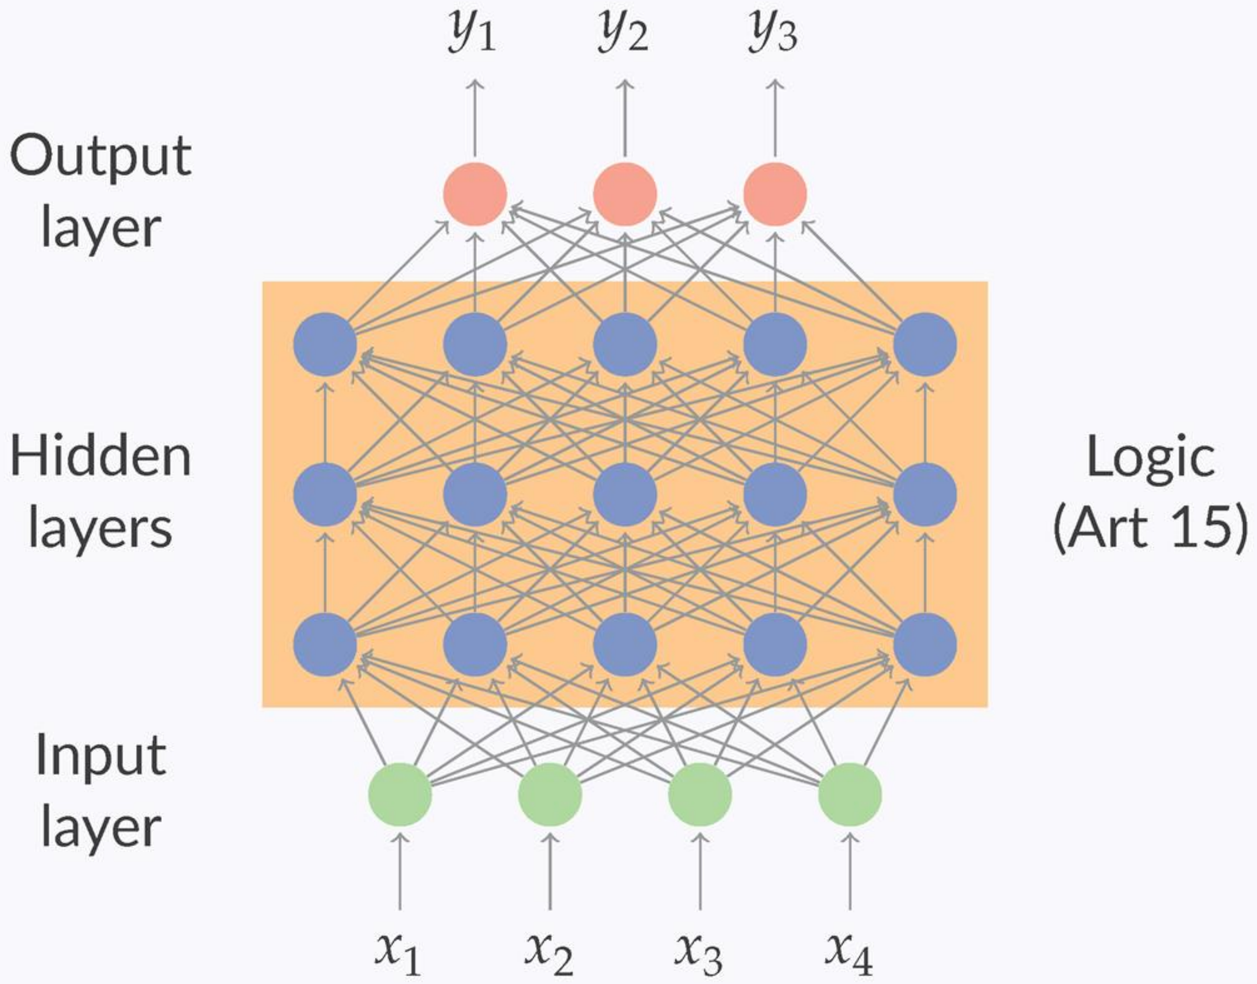

DEEP LEARNING AND THE GDPR

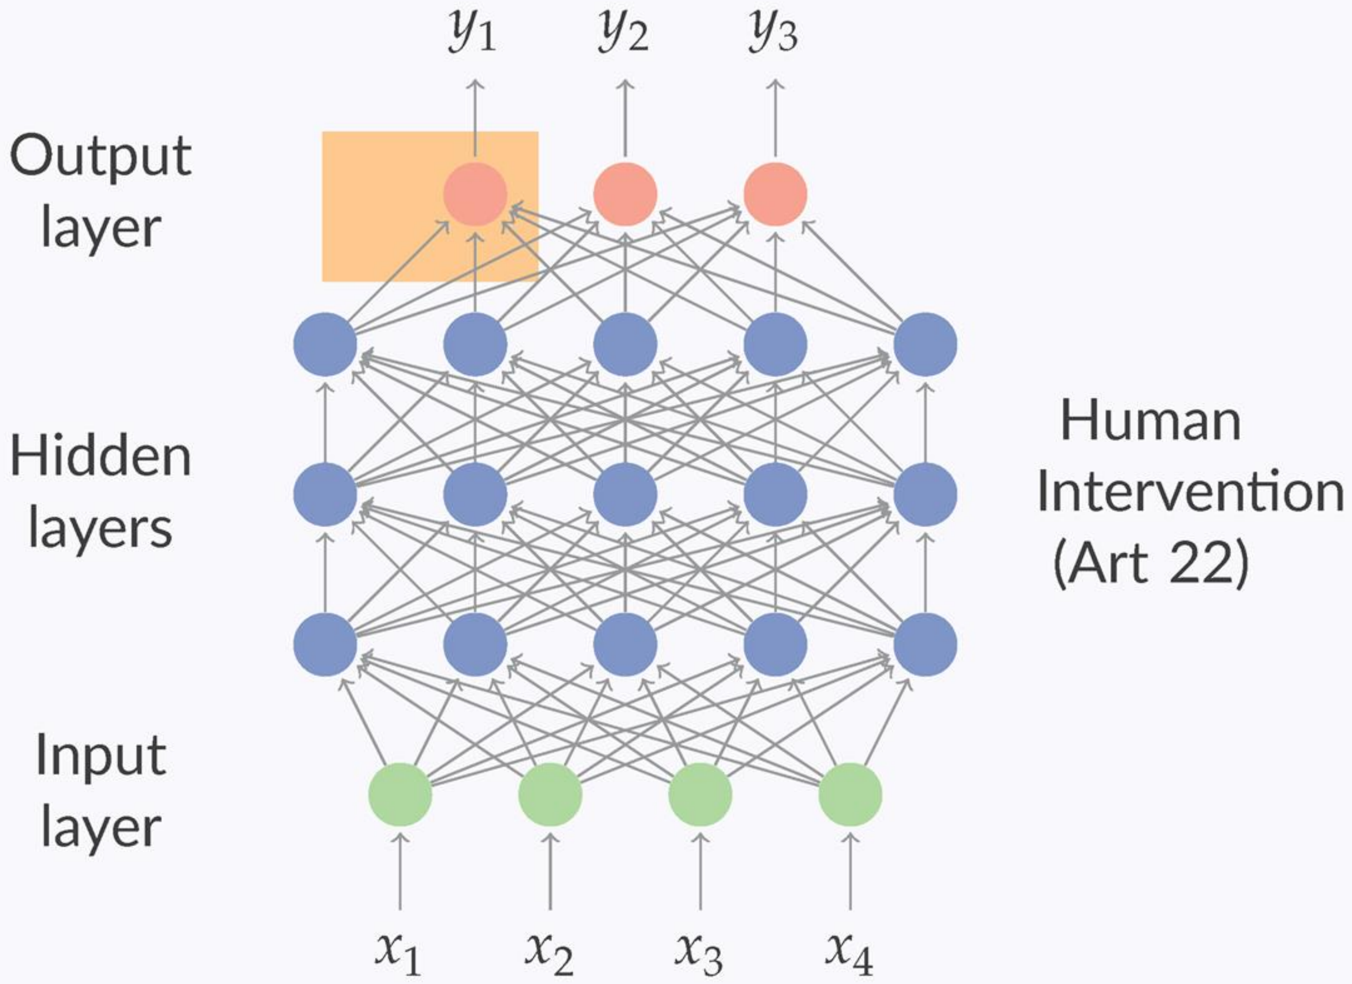

DEEP LEARNING AND THE GDPR

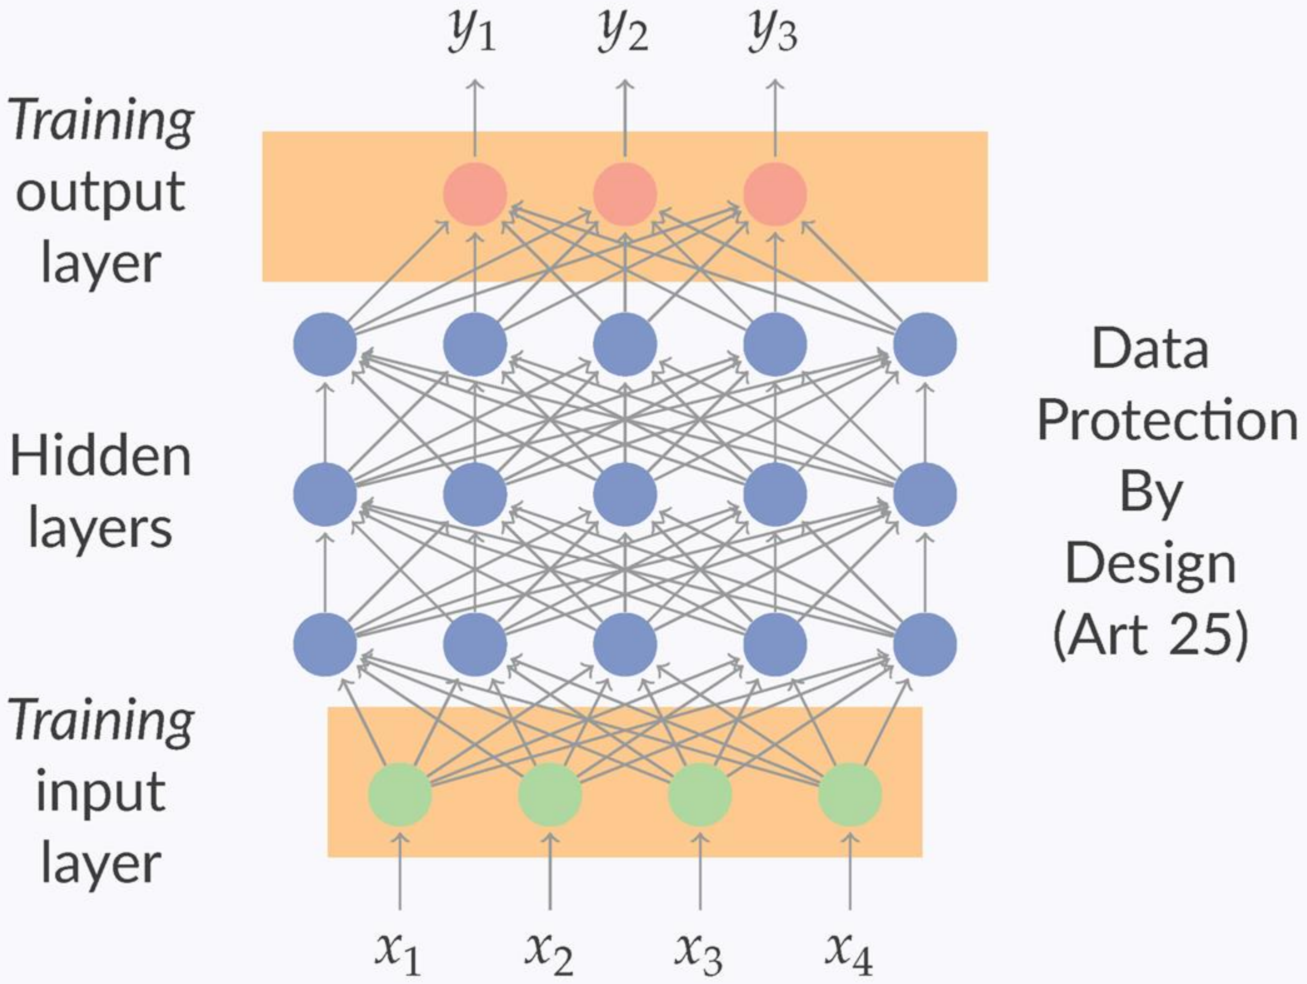

DEEP LEARNING AND THE GDPR

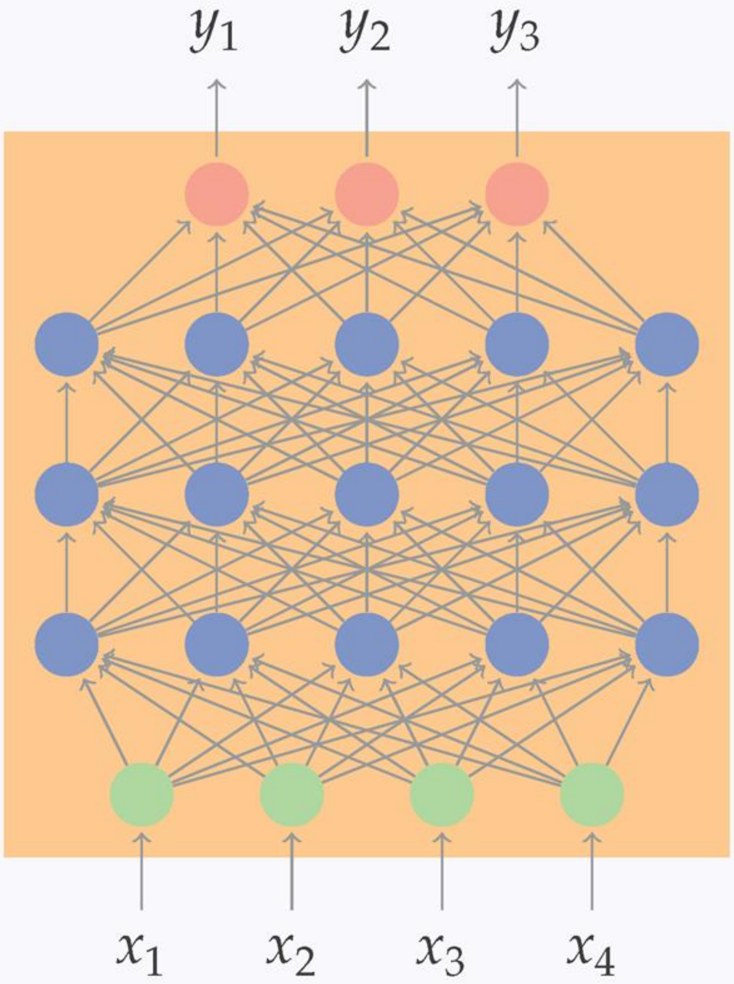

Training and real output layers

Hidden layers

Training and real input layers

Certification (Art 42)

QUESTIONS FOR A CERTIFICATION AUTHORITY

- What do we need from an authority?
- What information could be measured by an authority? Can harms/risks be measured or inferred?
- Should measurement be external (pedagogical/adversarial) or internal (an accountant/auditor)?
- Where do standards fit in?
- Certify the data, the algorithm, the data scientist or all of these?
- Who should be involved? IEEE?[7] BSI? Or a multistakeholder effort?

[7] The IEEE Global Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems. Ethically aligned design: A vision for prioritizing wellbeing with artificial intelligence and autonomous systems, 2016. Version 1, http://standards.ieee.org/develop/indconn/ec/autonomous_systems.html