ETH zürich

CPO – Regione Piemonte

L'intelligenza artificiale al servizio della salute: aspetti bioetici

Alessandro Blasimme, PhD

DHEST
Department of Health Sciences and Technology

Outline

1. From bodies to data
2. From doctors to (black-box) algorithms
3. Conclusions: towards responsible clinical use of AI

All of Us℠ | The Precision Medicine Initiative®

THE FUTURE OF HEALTH BEGINS WITH YOU

SPHN

Swiss Personalized Health Network

100,000 genomes

70,000 patients and family members

21 Petabytes of data.

1 Petabyte of music would take 2,000 years to play on an MP3 player.

Genomics england


Support The Guardian
Available for everyone, funded by readers
Contribute → Subscribe →
Search jobs | Sign in | Search | International edition
The Guardian

News | Opinion | Sport | Culture | Lifestyle | More

World UK Science Cities Global development Football **Tech** Business Environment Obituaries

Google

This article is more than **3 years old**

Google given access to healthcare data of up to 1.6 million patients

Artificial intelligence firm DeepMind provided with patient information as part of agreement with Royal Free NHS trust

Ben Quinn
@BenQuinn75
Wed 4 May 2016 00.34 BST

798 1,273

IBM Watson Health AI gets access to full health data of 61m Italians

Detailed medical records of 61 million Italian citizens to be given to IBM for its "cognitive computing" system Watson

Posted on May 22, 2017 by Glyn Moody

How care.data works

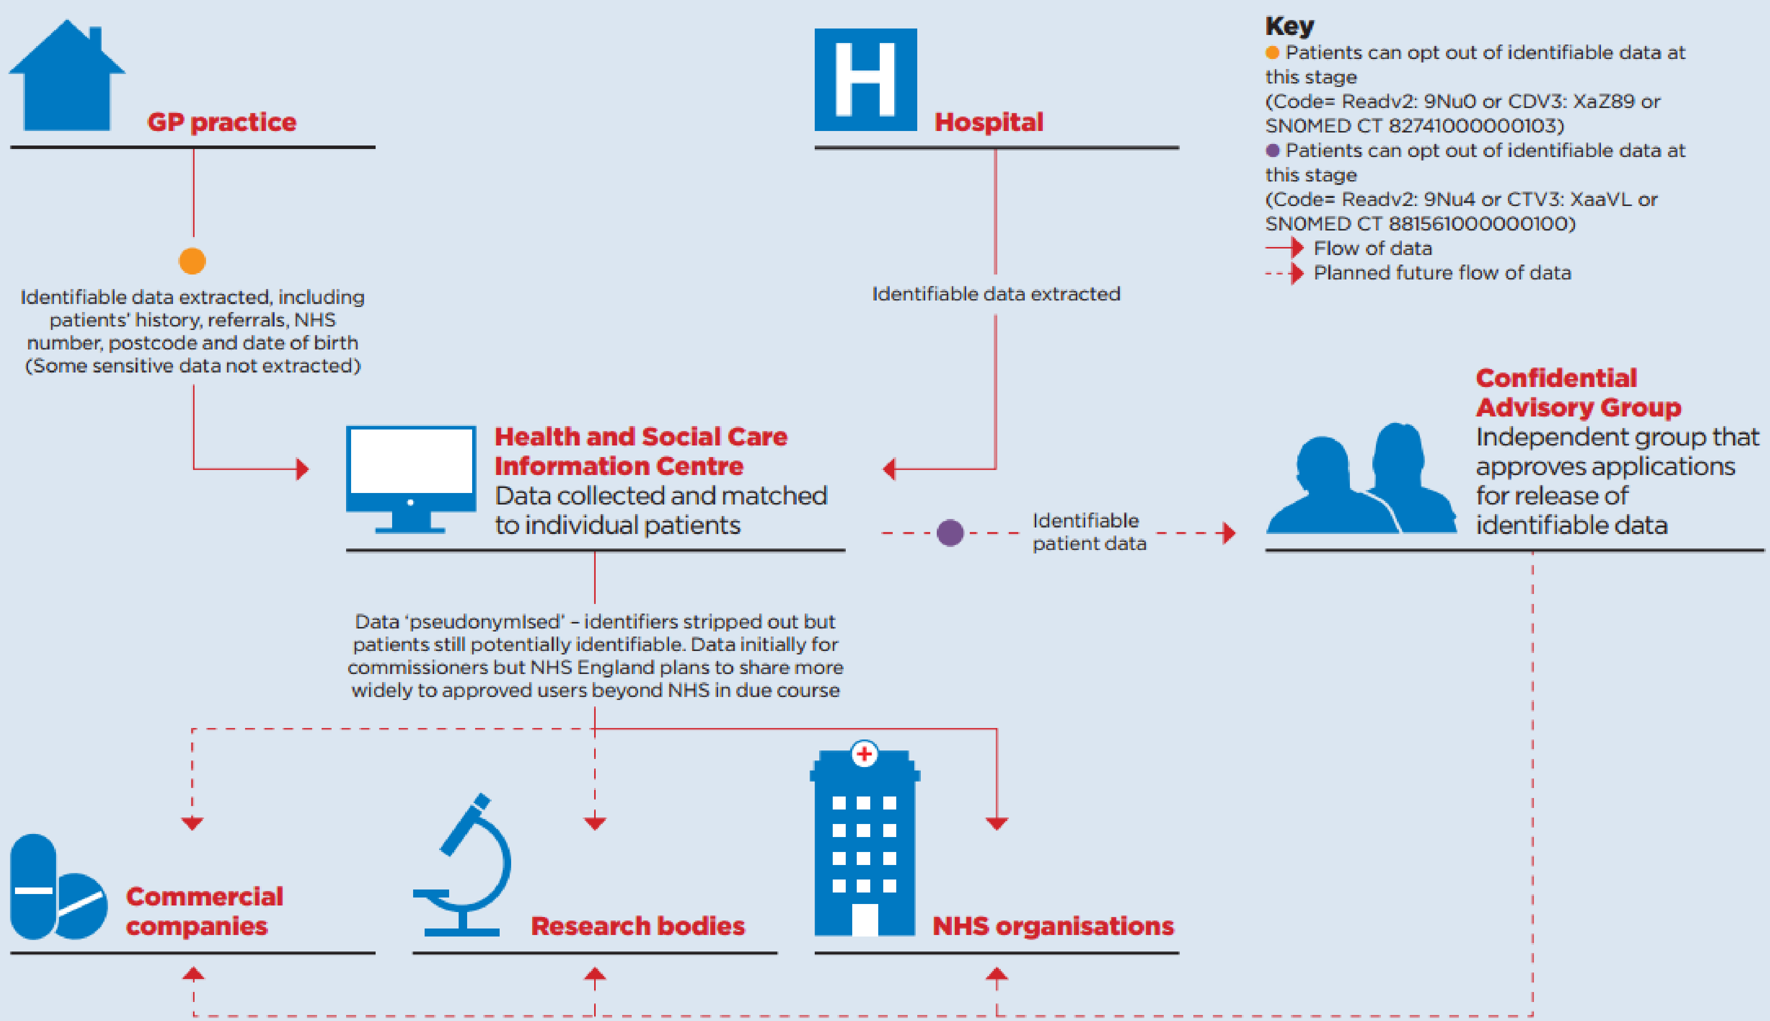

Support The Guardian

Available for everyone, funded by readers

Contribute → Subscribe →

Search jobs | Sign in | Search | International edition

The Guardian

News | Opinion | Sport | Culture | Lifestyle | More

World UK Science Cities Global development Football **Tech** Business Environment Obituaries

Data protection

This article is more than **3 years old**

NHS to scrap single database of patients' medical details

Care.data scheme to close after Fiona Caldicott review calls for tougher measures to keep information confidential

Sarah Boseley *Health editor*

Wed 6 Jul 2016 16.58 BST

Medical records

Electronic health records

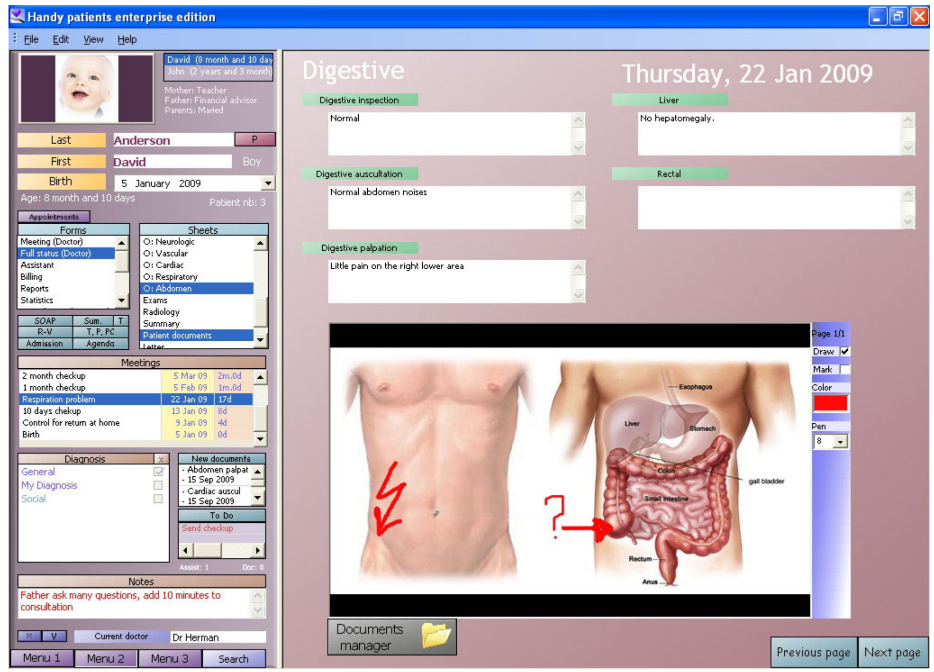

OPEN ACCESS Freely available online

PLoS COMPUTATIONAL BIOLOGY

Review

Digital Epidemiology

Marcel Salathé[1,2]*, Linus Bengtsson[3], Todd J. Bodnar[1,2], Devon D. Brewer[4], John S. Brownstein[5], Caroline Buckee[6], Ellsworth M. Campbell[1,2], Ciro Cattuto[7], Shashank Khandelwal[1,2], Patricia L. Mabry[8], Alessandro Vespignani[9]

1 Center for Infectious Disease Dynamics, Penn State University, University Park, Pennsylvania, United States of America, **2** Department of Biology, Penn State University, University Park, Pennsylvania, United States of America, **3** Department of Public Health Sciences, Karolinska Institutet, Stockholm, Sweden, **4** Interdisciplinary Scientific Research, Seattle, Washington, United States of America, **5** Harvard Medical School and Children's Hospital Informatics Program, Boston, Massachusetts, United States of America, **6** Center for Communicable Disease Dynamics, Department of Epidemiology, Harvard School of Public Health, Boston, Massachusetts, United States of America, **7** Institute for Scientific Interchange (ISI) Foundation, Torino, Italy, **8** Office of Behavioral and Social Sciences Research, NIH, Bethesda, Maryland, United States of America, **9** College of Computer and Information Sciences and Bouvé College of Health Sciences, Northeastern University, Boston, Massachusetts, United States of America

Abstract: Mobile, social, real-time: the ongoing revolution in the way people communicate has given rise to a new kind of epidemiology. Digital data sources, when harnessed appropriately, can provide local and timely information about disease and health dynamics in populations around the world. The rapid, unprecedented increase in the availability of relevant data from various digital sources creates considerable technical and computational challenges.

npj Digital Medicine
www.nature.com/npjdigitalmed

PERSPECTIVE OPEN

Data mining for health: staking out the ethical territory of digital phenotyping

Nicole Martinez-Martin[1], Thomas R. Insel[2], Paul Dagum[2], Henry T. Greely[1] and Mildred K. Cho[1]

PLOS | MEDICINE

PERSPECTIVE

Machine learning in medicine: Addressing ethical challenges

Effy Vayena[1*], Alessandro Blasimme[1], I. Glenn Cohen[2]

1 Health Ethics and Policy Lab, Department of Health Sciences and Technology, ETH Zurich, Zurich, Switzerland, 2 Harvard Law School, Cambridge, Massachusetts, United States of America

Review article: Biomedical intellegence | Published 16 January 2018 | doi:10.4414/smw.2018.14571
Cite this as: Swiss Med Wkly. 2018;148:w14571

Digital health: meeting the ethical and policy challenges

Vayena Effy[a], Haeusermann Tobias[b], Adjekum Afua[a], Blasimme Alessandro[a]

[a] Health Ethics and Policy Lab, Department of Health Sciences and Technology (D-HEST), ETH Zurich, Switzerland
[b] Department of Sociology, University of Cambridge, UK

N Engl J Med. 2018 March 15; 378(11): 981–983. doi:10.1056/NEJMp1714229.

Implementing Machine Learning in Health Care — Addressing Ethical Challenges

Danton S. Char, M.D., **Nigam H. Shah, M.B., B.S., Ph.D.**, and **David Magnus, Ph.D.**

The Ethics of AI in Biomedical Research, Patient Care and Public Health

Oxford Handbook of Ethics of Artificial Intelligence, Forthcoming

25 Pages • Posted: 17 May 2019

Alessandro Blasimme
ETH Zurich

Effy Vayena
ETH Zurich

https://doi.org/10.2139/ssrn.3368756

AMA Journal of Ethics®

AMA

FEBRUARY 2019

Artificial Intelligence in Health Care

In health care, artificial intelligence (AI) can help manage and analyze data, make decisions, and conduct conversations, so it is destined to drastically change clinicians' roles and everyday practices. Adaptability to change in diagnostics, therapeutics, and practices of maintaining patients' safety and privacy will be key. This issue also explores some of the most ethically complex questions about AI's implementation, uses, and limitations in health care.

Volume 21, Number 2: E119-197
Full Issue PDF

The rise of digital health: ethical issues

- Data ethics
- Informed consent
- Explainability
- Patient-doctor relation
- Trust
- Burden of care
- Sustainability
- Medicalization
- Privacy
- Data protection
- Discrimination
- Evidence and validation
- Regulatory standards
- Liability

The social trading of personal data

- Under which conditions are personal data made available?
- What forms of individual and collective control exist?
- Are appropriate safeguards in place?
- Who benefits the most?
- How are decisions being taken?

The social trading of personal data: available solutions

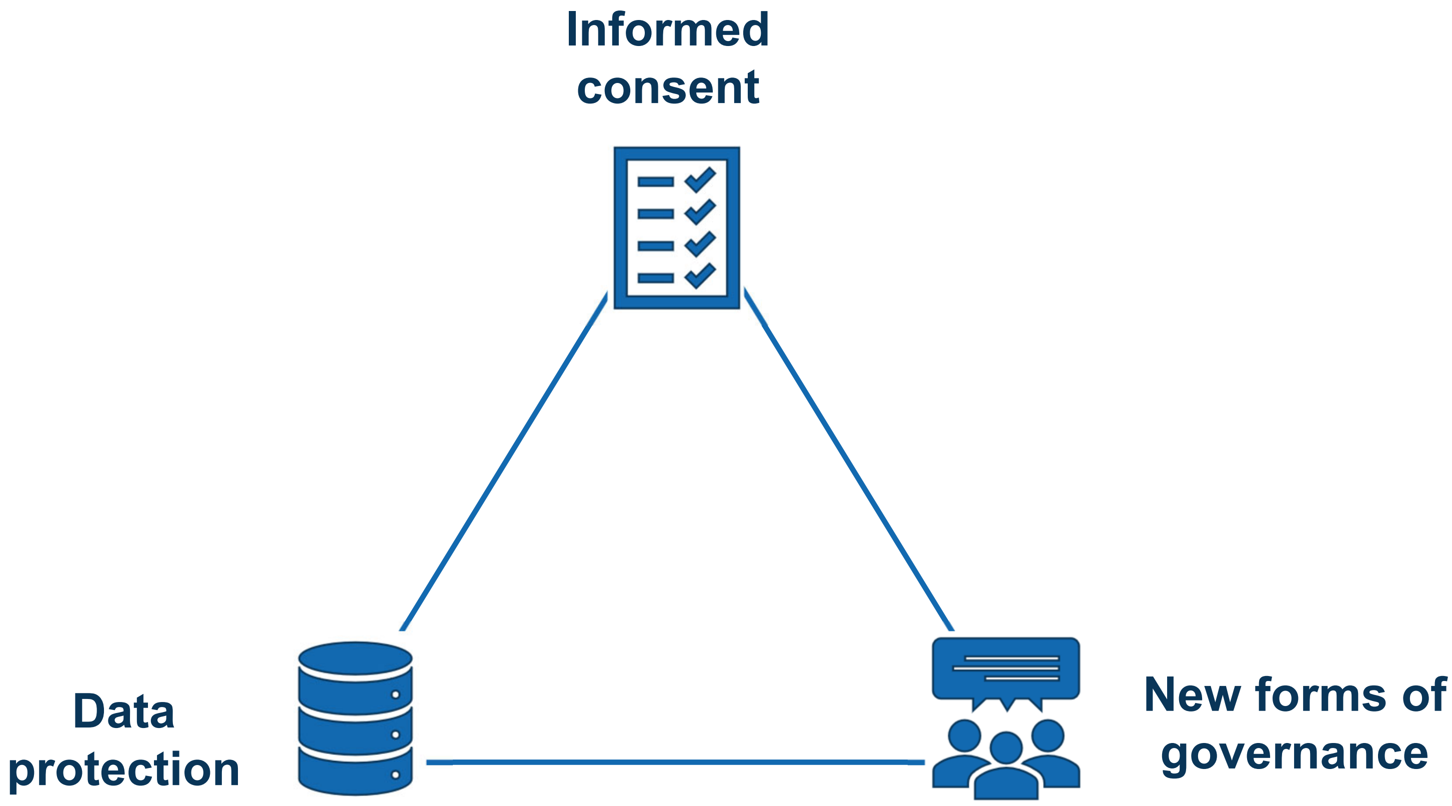

Informed consent

- Autonomous authorization to use data: agree to specific conditions of exposure
- Only limited amounts of control over the production, collection, use, and circulation of health data.
- Not a sufficient condition to ensure protection against privacy-related harms: e.g. discrimination, stigmatization, unfairness

A BROKEN CONTRACT

ate in May, the direct-to-consumer
ting company 23andMe proudly
ced the impending award of
patent. The firm's research on
sease, which used data from sev-
customers, had led to a patent on
es that contribute to risk for the
ight be used to predict its course.
ki, co-founder of the company,
w, California,
nt would help
n of academic

AS RESEARCHERS FIND MORE USES FOR DATA, INFORMED CONSENT HAS BECOME A SOURCE OF CONFUSION. SOMETHING HAS TO CHANGE.

BY ERIKA CHECK HAYDEN

ELSEVIER

The American Journal of the Medical Sciences

Volume 342, Issue 4, October 2011, Pages 267-272

Symposium Article

Is Informed Consent Broken?

Gail E. Henderson PhD

Renovating consent

- Info - control + data sharing	Midway		+ info +control - data sharing
	Overseen	Choice-based	
No consent **Presumed consent** (Gill 2004) **Presumed consent with opt-out** (Wendler and Emanuel 2002) **Blanket consent** (UNESCO 2001) (Tomlinson 2013) **Open consent** (Lunshof et al. 2008) **Portable legal consent** (Hayden 2012; Vayena, Mastroianni, and Kahn 2013) **Broad consent 1** = blanket + limitations (as defined in Grady et al. 2015)	**Broad consent 2** = blanket consent + safety + withdrawal + access review (Hansson et al. 2006) **Broad consent + ongoing oversight and communication** (Grady et al. 2015) **Broad consent + governance** (O'Doherty et al. 2011) **Broad consent + trusted governance system** (Koenig 2014; Garrett, Dohan, and Koenig 2015)	**Authorization model** (Caulfield, Upshur, and Daar 2003) **Tiered consent** (McGuire and Beskow 2010; Mello and Wolf 2010; Bunnik, Janssens, and Schermer 2013) **Electronic informed Consent** (FDA and DHHS 2016; Sage Bionetworks 2017) **Dynamic consent** (Kaye et al. 2012; Kaye et al. 2015; Budin-Ljøsne et al. 2017)	**Informed consent** (Faden and Beauchamp 1986; Manson and O'Neill 2007) **Consent for de-identified samples and data**

Data protection

- Pseudonymization
- Anonymization
- Encryption

The more data circulate, the harder it is to protect it.

Identifying Personal Genomes by Surname Inference

Melissa Gymrek,[1,2,3,4] Amy L. McGuire,[5] David Golan,[6] Eran Halperin,[7,8,9] Yaniv Erlich[1*]

Sharing sequencing data sets without identifiers has become a common practice in genomics. Here, we report that surnames can be recovered from personal genomes by profiling short tandem repeats on the Y chromosome (Y-STRs) and querying recreational genetic genealogy databases. We show that a combination of a surname with other types of metadata, such as age and state, can be used to triangulate the identity of the target. A key feature of this technique is that it entirely relies on free, publicly accessible Internet resources. We quantitatively analyze the probability of identification for U.S. males. We further demonstrate the feasibility of this technique by tracing back with high probability the identities of multiple participants in public sequencing projects.

SCIENCE VOL 339 18 JANUARY 2013

Identification of individuals by trait prediction using whole-genome sequencing data

Christoph Lippert[a,1], Riccardo Sabatini[a], M. Cyrus Maher[a], Eun Yong Kang[a], Seunghak Lee[a], Okan Arikan[a], Alena Harley[a], Axel Bernal[a], Peter Garst[a], Victor Lavrenko[a], Ken Yocum[a], Theodore Wong[a], Mingfu Zhu[a], Wen-Yun Yang[a], Chris Chang[a], Tim Lu[b], Charlie W. H. Lee[b], Barry Hicks[a], Smriti Ramakrishnan[a], Haibao Tang[a], Chao Xie[c], Jason Piper[c], Suzanne Brewerton[c], Yaron Turpaz[b,c], Amalio Telenti[b], Rhonda K. Roby[b,d,2], Franz J. Och[a], and J. Craig Venter[b,d,1]

10166–10171 | PNAS | September 19, 2017 | vol. 114 | no. 38

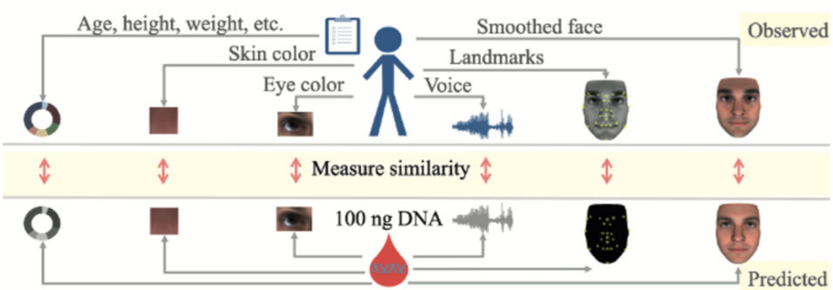

Fig. 6. Overview of the experimental approach. A DNA sample and a vari-

New forms of governance

- Inclusive governance of research data:
  - Publicly sponsored research

All of Us℠
THE FUTURE OF HEALTH BEGINS WITH YOU
The Precision Medicine Initiative®

"TAILORED-TO-YOU"

public engagement and the political legitimation of precision medicine

https://muse.jhu.edu/article/648044/pdf

ALESSANDRO BLASIMME AND EFFY VAYENA

Blasimme and Vayena *BMC Medical Ethics* (2016) 17:67
DOI 10.1186/s12910-016-0149-6

BMC Medical Ethics

DEBATE Open Access

Becoming partners, retaining autonomy: ethical considerations on the development of precision medicine

Alessandro Blasimme[1,2*] and Effy Vayena[1]

New forms of governance

- Inclusive governance of research data:
  - Publicly sponsored research
  - Data cooperatives

Philos. Technol. (2018) 31:473–479
https://doi.org/10.1007/s13347-018-0320-8

COMMENTARY

Democratizing Health Research Through Data Cooperatives

Alessandro Blasimme[1] · Effy Vayena[1] · Ernst Hafen[2]

Outline

1. From bodies to data
2. From doctors to (black-box) algorithms
3. Conclusions: towards responsible clinical use of AI

AI in society

For artificial intelligence to thrive, it must explain itself

If it cannot, who will trust it?

Print edition | Science and technology Feb 15th 2018

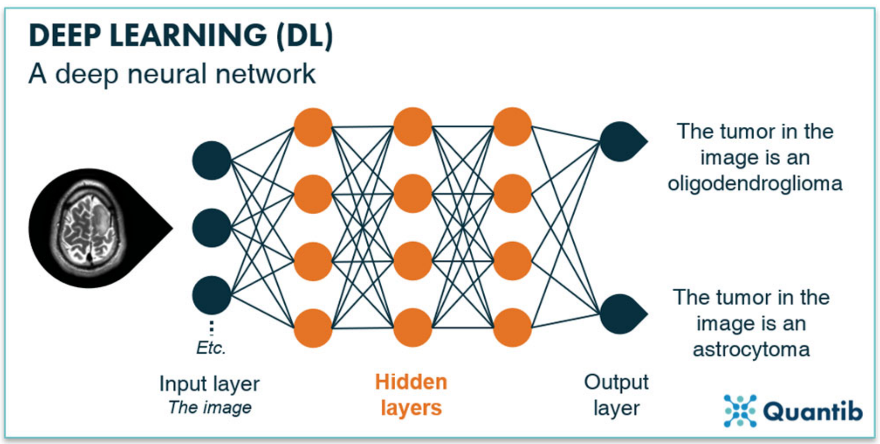

SCIENCE ROBOTICS | FOCUS

ROBOTS AND SOCIETY

Transparent, explainable, and accountable AI for robotics

Sandra Wachter,[1]* Brent Mittelstadt,[2,3,1] Luciano Floridi[1,2]

To create fair and accountable AI and robotics, we need precise regulation and better methods to certify, explain, and audit inscrutable systems.

International Data Privacy Law, 2017, Vol. 7, No. 4 ARTICLE

Meaningful information and the right to explanation

Andrew D. Selbst* and Julia Powles**

SCIENCE TRANSLATIONAL MEDICINE | FOCUS

BIG DATA

Big data and black-box medical algorithms

W. Nicholson Price[1,2,3]

New machine-learning techniques entering medicine present challenges in validation, regulation, and integration into practice.

Artificial Intelligence and Black-Box Medical Decisions: *Accuracy versus Explainability*

THE HASTINGS CENTER REPORT

BY ALEX JOHN LONDON

A right to explanation in the GDPR?

- Recital 71
- Articles 13, 14, 15
- Data subjects are entitled to receive meaningful information about the logic involved, the significance and the envisaged consequences of solely automated individual decision-making and profiling.

Machine Learning in Medicine:

Opening the New Data Protection Black Box

*Agata Ferretti, Manuel Schneider and Alessandro Blasimme**

DOI: 10.21552/edpl/2018/3/10

What does *opacity* even mean?

Machine Learning in Medicine:
Opening the New Data Protection Black Box
*Agata Ferretti, Manuel Schneider and Alessandro Blasimme**

- OPACITY 1:
  - It is impossible to access the rules the algorithm has learnt: not programmed.
  - It is impossible to make sense of the rules the algorithm applies: too complex

- OPACITY 2:
  - It is impossible to understand why AI makes this or that decision/prediction: why are input and output are associated?

Why must AI be explainable?

- Clinical interpretability:
  - Why did the AI do that?
  - Why did the AI succeed?
  - Why did it fail?
- To correct AI's mistake
- To ensure AI can be trustworthy

Shedding light on darkness

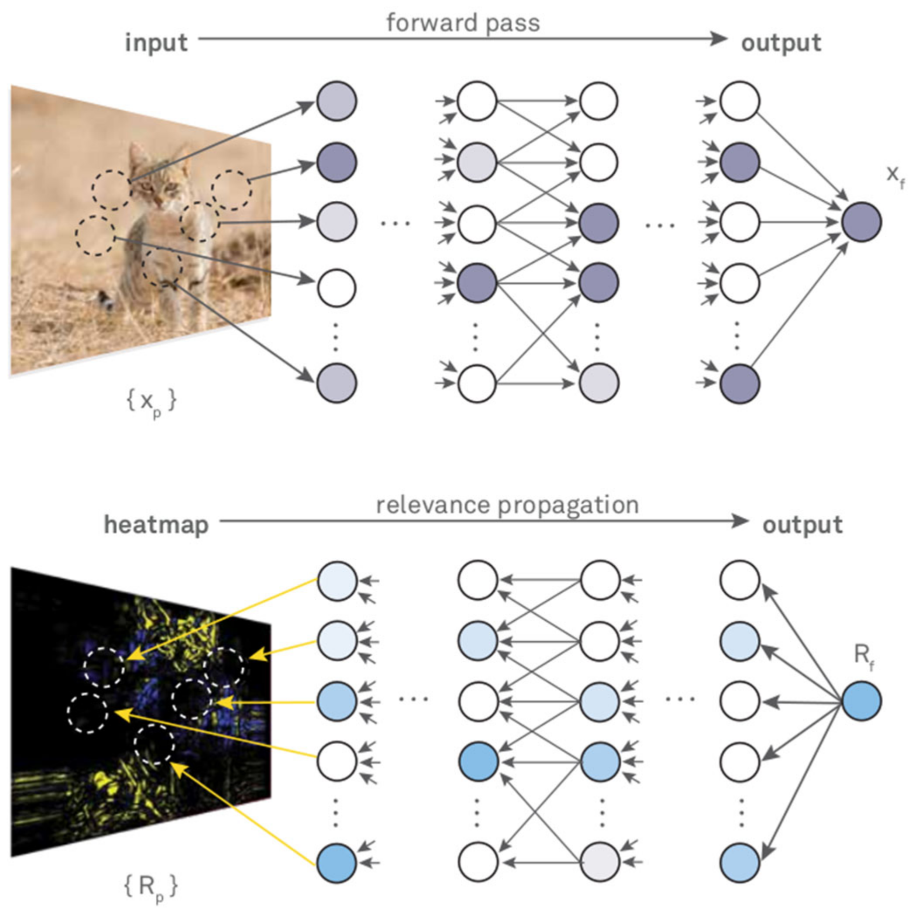

Wipro.com

The heatmap provides visual depiction of the relevance of a feature in the decision making. In an image classifier, it represents the contribution of the pixels towards a class.

Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning

nature biomedical engineering
https://doi.org/10.1038/s41551-018-0195-0

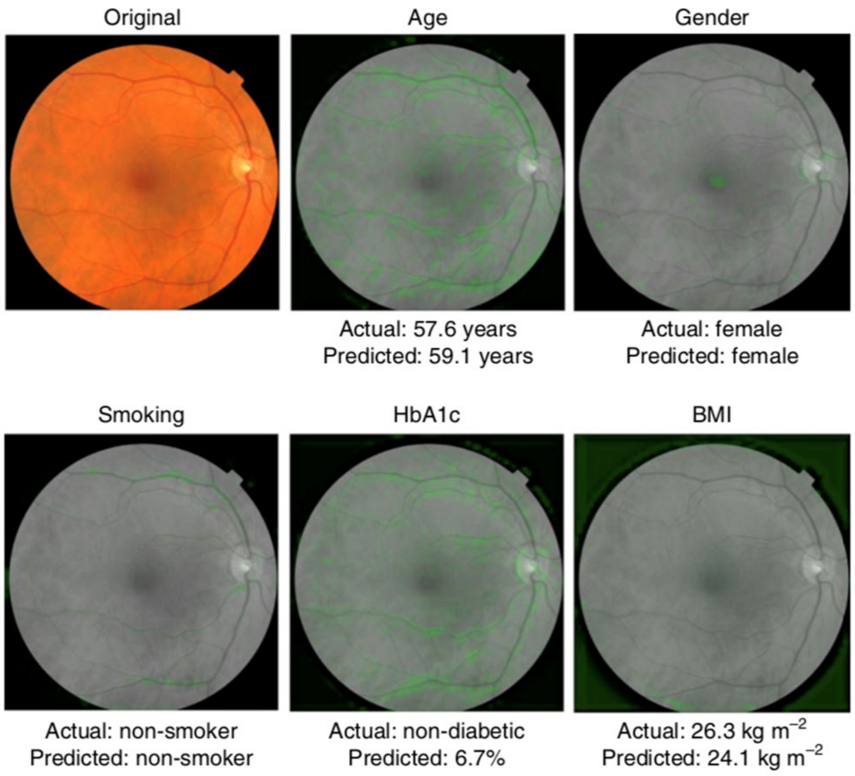

Fig. 2 | Attention maps for a single retinal fundus image. The top left image is a sample retinal image in colour from the UK Biobank dataset. The remaining images show the same retinal image, but in black and white. The soft attention heat map (Methods) for each prediction is overlaid in green, indicating the areas of the heat map that the neural-network model is using to make the prediction for the image. For a quantitative analysis of what was highlighted, see Table 6. HbA1c values are not available for UK Biobank patients, so the self-reported diabetes status is shown instead.

Opening the Black Box of Artificial Intelligence for Clinical Decision Support: A Study Predicting Stroke Outcome

Esra Zihni[1*], Vince Istvan Madai[1*], Michelle Livne[1], Ivana Galinovic[2], Ahmed A. Khalil[2], Jochen B. Fiebach[2], Dietmar Frey[1]

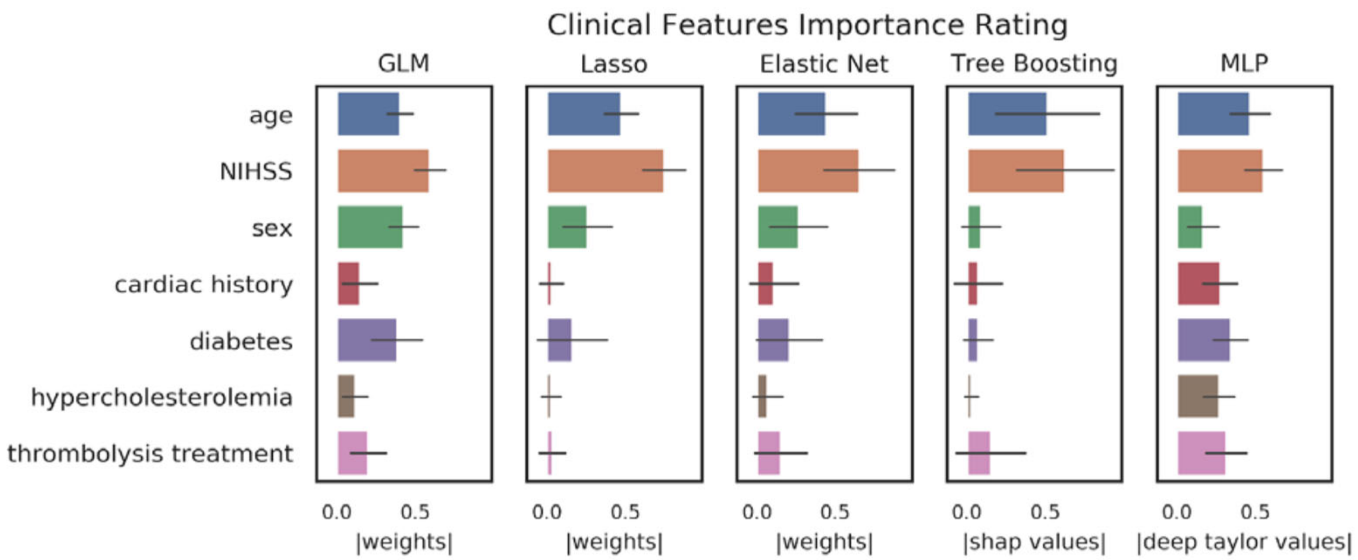

Fig 3. Graphical representation of the feature importance. The figure illustrates the features rating derived from the model-tailored interpretability methods for generalized linear model (GLM), Lasso, Elastic net, Catboost and multilayer perceptron (MLP). For logistic regression techniques the results are given in weights, for Catboost in Shap(ley) values and for MLP in deep Taylor values that were normalized to the range [0,1]. The bar heights represent means and error bars represent standard deviation over samples (shuffles).

Is correlation without *explanation* acceptable?

- Can we accept to perform a medical act we don't know how to explain?

- *Explanatory* reasons vs *Justificatory* reasons
- AI will never provide explanatory reasons in mechanistic terms: only the lab can do that.
- Mechanisms are neither necessary nor sufficient as justificatory reasons for a medical act.

Is correlation without explanation useful?

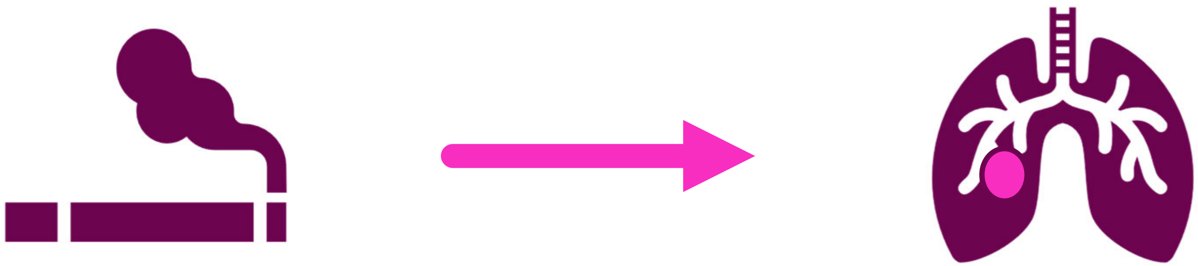

Discovery of causes (causal inference)

Is correlation without explanation useful?

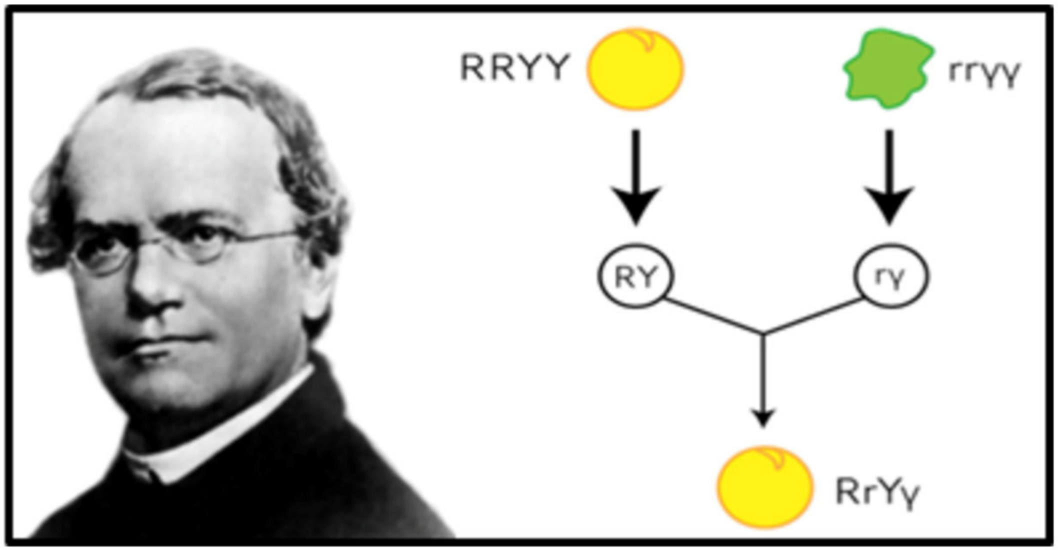

Gregor Mendel, F1 dihybrid; © CC AndreaLaurel

Generation of new hypotheses

Does correlation without *explanation* constitute evidence?

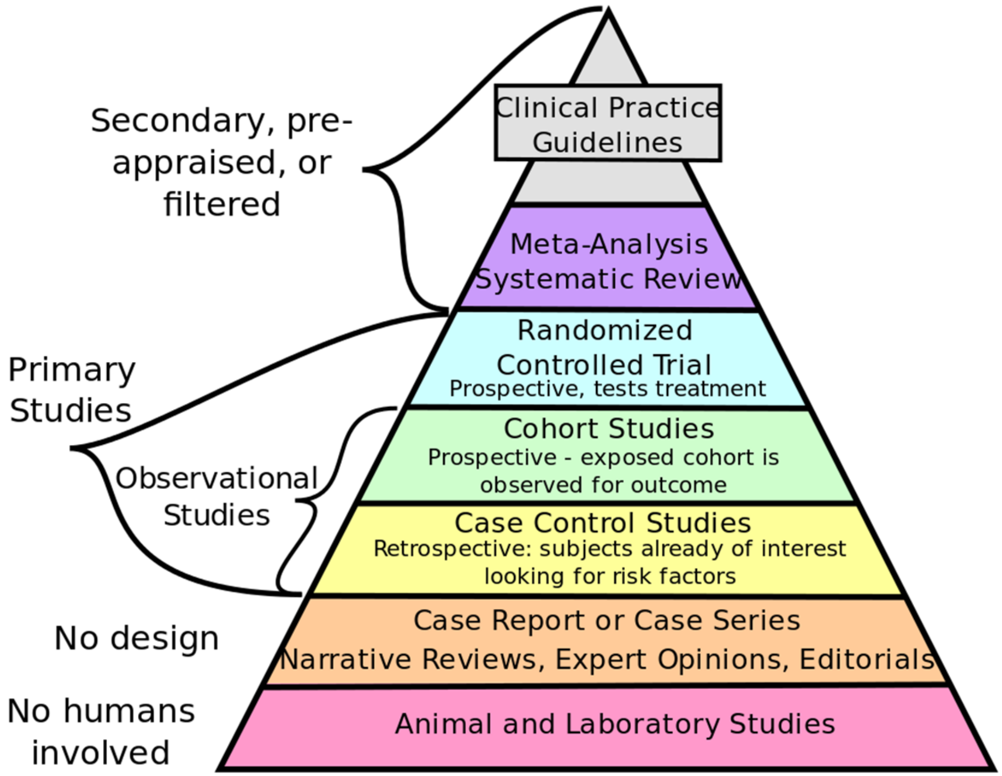

J Med Internet Res. 2019 Mar; 21(3): e10769.

Probabilistic evidence

Evidence and explanation in medicine

- Multiple sources (and formats) of medical evidence: probabilistic + mechanistic.

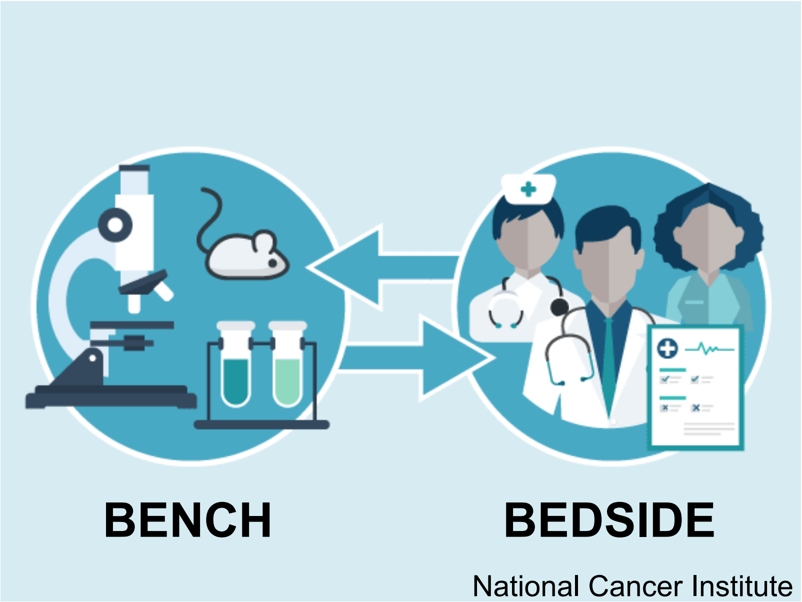

Medicine relies on **evidence integration**

(Russo & Williamson 2007)

Protecting evidence integration

- AI should not get in the way of evidence integration:
  - Disproportionate R&D funding (undermining basic research)
  - Economic incentives to clinical use
  - Blind faith / trust
  - Clinical dependency
  - Deskilling

What kind of explanations do *we* need?

- As patients…
- As health care professionals…
- As scientists…

Outline

1. From bodies to data
2. From doctors to (black-box) algorithms
3. Conclusions: towards responsible clinical use of AI

Three maxims for a technological world

1. That technologies be given a scale and structure of the sort that would be immediately intelligible to nonexperts.
2. That technologies be built with a high degree of flexibility and mutability.
3. That technologies be judged according to the degree of dependence they tend to foster, those creating the greater dependency being held inferior.

L. Winner, 1977, p. 326-7

Grazie per l'attenzione!

Health Ethics
& Policy Lab

alessandro.blasimme@hest.ethz.ch

Twitter: **a_blasimme**