MACQUARIE University

The impact of AI on medicine and medical education

@enricocoiera

The New York Times Magazine

The Great A.I. Awakening

How Google used artificial intelligence to transform Google Translate, one of its more popular services — and how machine learning is poised to reinvent computing itself.

BY GIDEON LEWIS-KRAUS DEC. 14, 2016

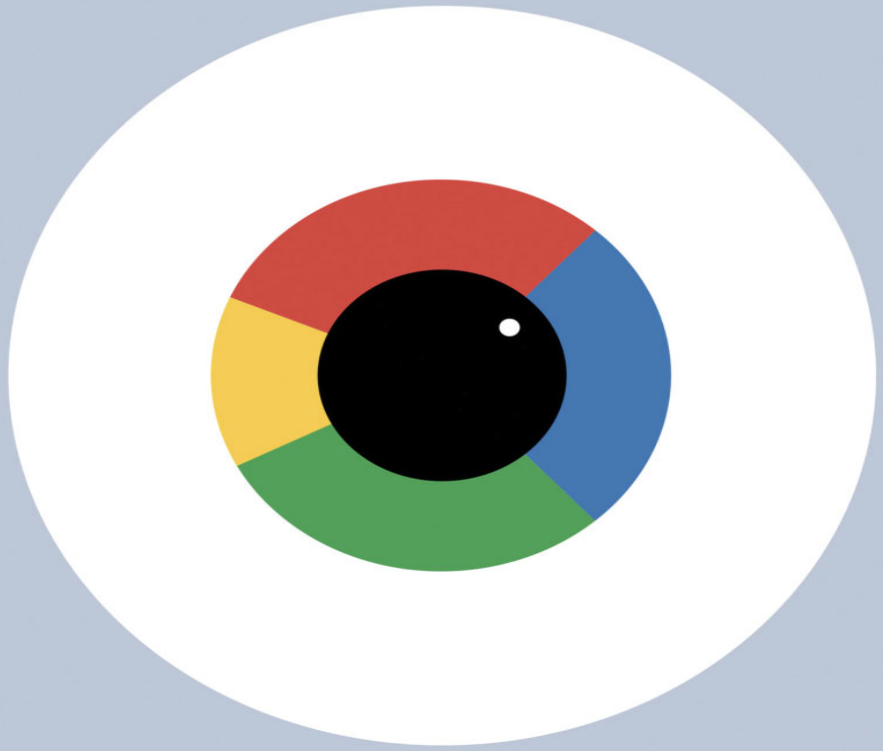

radiology.rsna.org ▪ **Radiology:** Volume 284: Number 2—August 2017

Radiology

Deep Learning at Chest Radiography: Automated Classification of Pulmonary Tuberculosis by Using Convolutional Neural Networks[1]

Results: The best-performing classifier had an AUC of 0.99, which was an ensemble of the AlexNet and GoogLeNet DCNNs. The AUCs of the pretrained models were greater than that of the untrained models ($P < .001$). Augmenting the dataset further increased accuracy (P values for AlexNet and GoogLeNet were .03 and .02, respectively). The DCNNs had disagreement in 13 of the 150 test cases, which were blindly reviewed by a cardiothoracic radiologist, who correctly interpreted all 13 cases (100%). This radiologist-augmented approach resulted in a sensitivity of 97.3% and specificity 100%.

MIT Technology Review

Intelligent Machines

IBM's Automated Radiologist Can Read Images and Medical Records

Software that can read medical images and written health records could help radiologists work faster and more accurately.

by Tom Simonite February 4, 2016

Most smart software in use today specializes on one type of data, be that interpreting text or guessing at the content of photos. Software in development at IBM has to do all those at once. It's in training to become a radiologist's assistant.

ANNALS OF MEDICINE APRIL 3, 2017 ISSUE

A.I. VERSUS M.D.

What happens when diagnosis is automated?

By Siddhartha Mukherjee

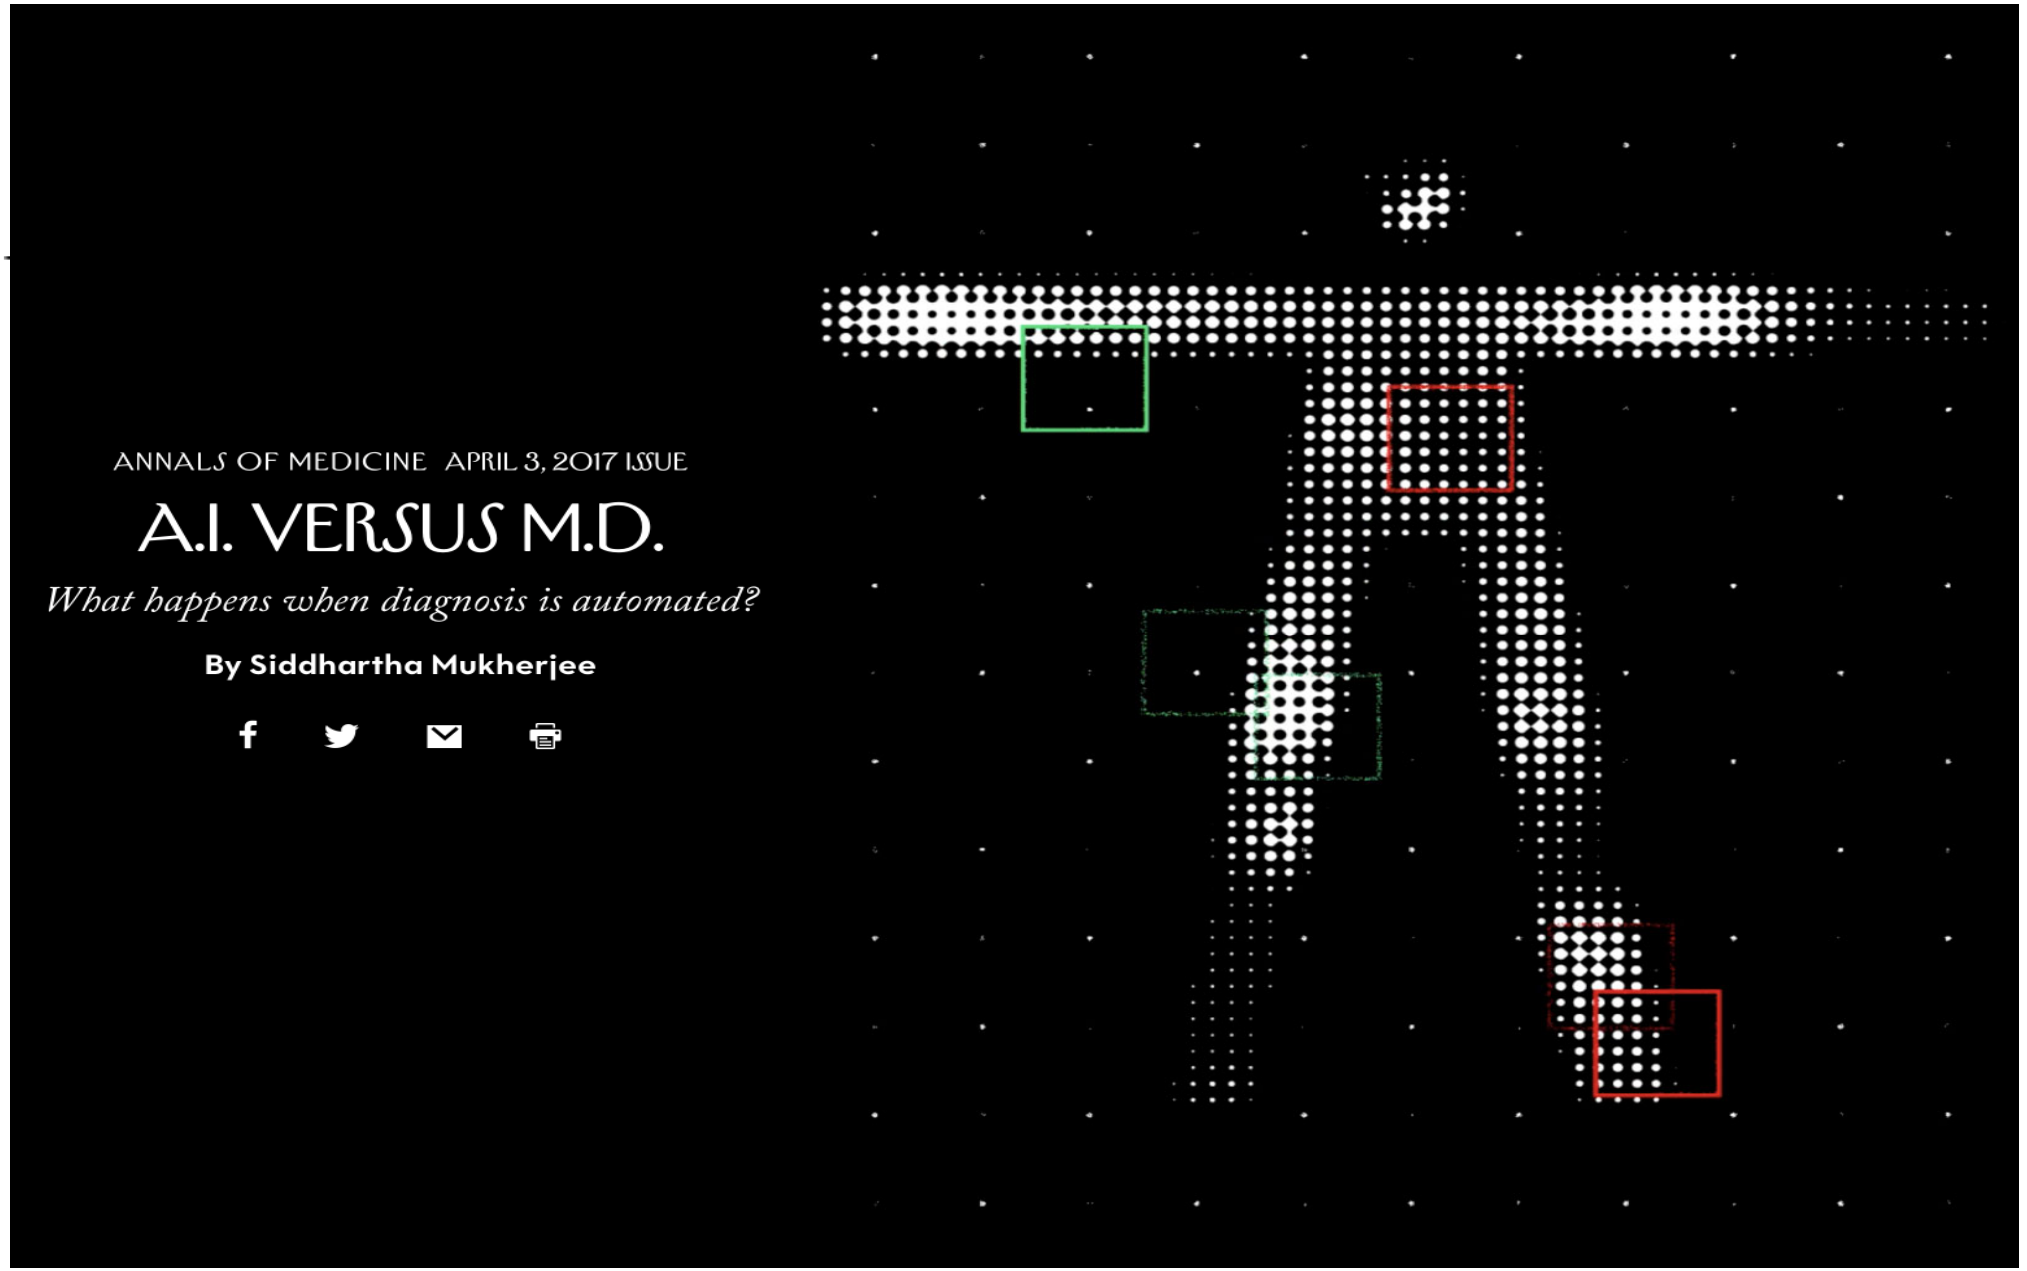

Perspective

COMMENTARY

Will Computers Replace Radiologists?

Saurabh Jha, MBBS, MRCS

DISCLOSURES | May 12, 2016

I recently told a radiology resident who demolished the worklist, "You're a machine." He beamed with pride. Imitation is the highest form of flattery. But the machine, not content in being imitated, wants to encroach on our turf.

CT could scarcely have progressed without progress in computing to horde the glut of thin slices. On two-dimensional projectional images such as abdominal radiographs, arteries reveal themselves only if calcified. With CT, algorithms extract arteries with a click. When I first saw this I was mesmerized, then humbled, and then depressed. With automation, what was my role, aside from clicking the aorta?

The role of computers in radiology was predicted as early as 1959 by Lee Lusted,[1] a radiologist and the founder of the Society for Medical Decision Making. Lusted envisioned "an electronic scanner-computer to look at chest photofluorograms, to separate the clearly normal chest films from the abnormal chest films. The abnormal chest films would be marked for later study by the radiologists."

AI Is Not Just a Sci-Fi Film

MEGAN MOLTENI SCIENCE 01.25.17 01:00 PM

IF YOU LOOK AT X-RAYS OR MOLES FOR A LIVING, AI IS COMING FOR YOUR JOB

Australian Alliance
for Artificial Intelligence in Healthcare

WIRED

THE WALL STREET JOURNAL.

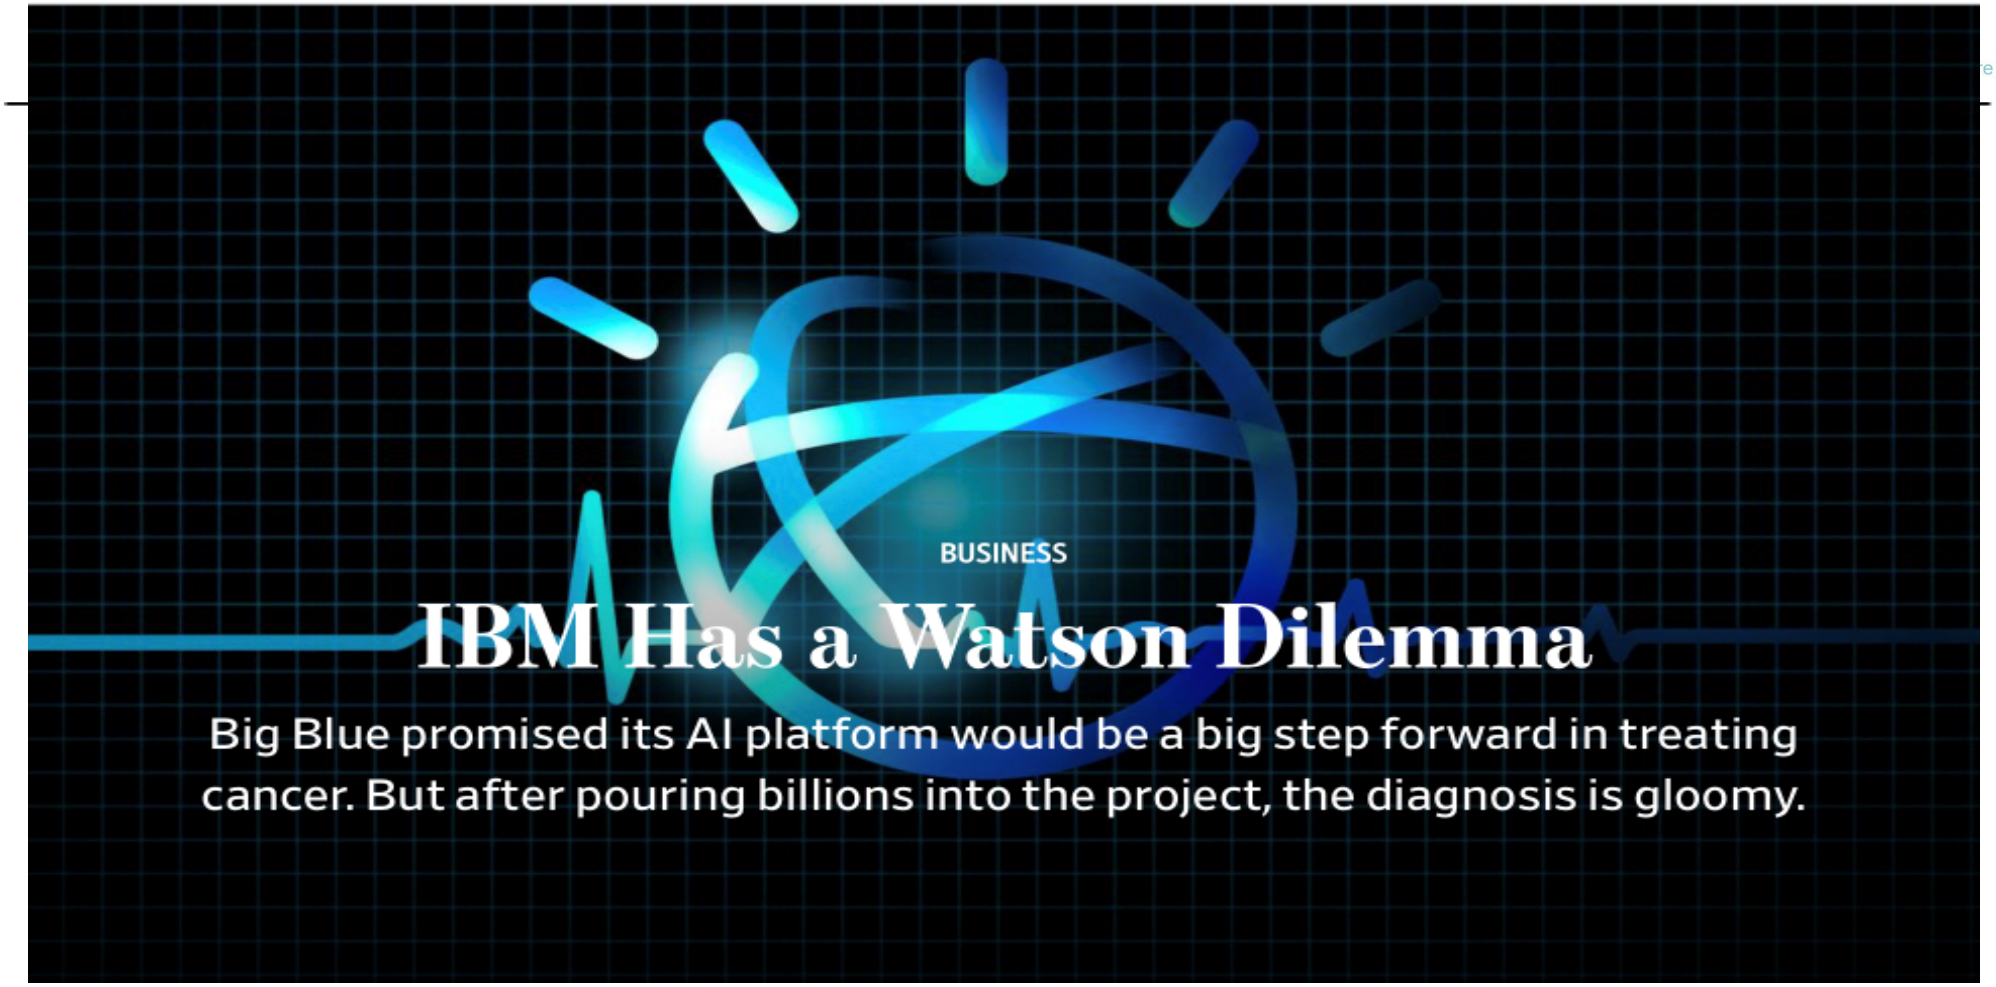

BUSINESS

IBM Has a Watson Dilemma

Big Blue promised its AI platform would be a big step forward in treating cancer. But after pouring billions into the project, the diagnosis is gloomy.

FINANCIAL TIMES

High-profile health app under scrutiny after doctors' complaints

Babylon advice service faces warnings it can miss symptoms of serious illness

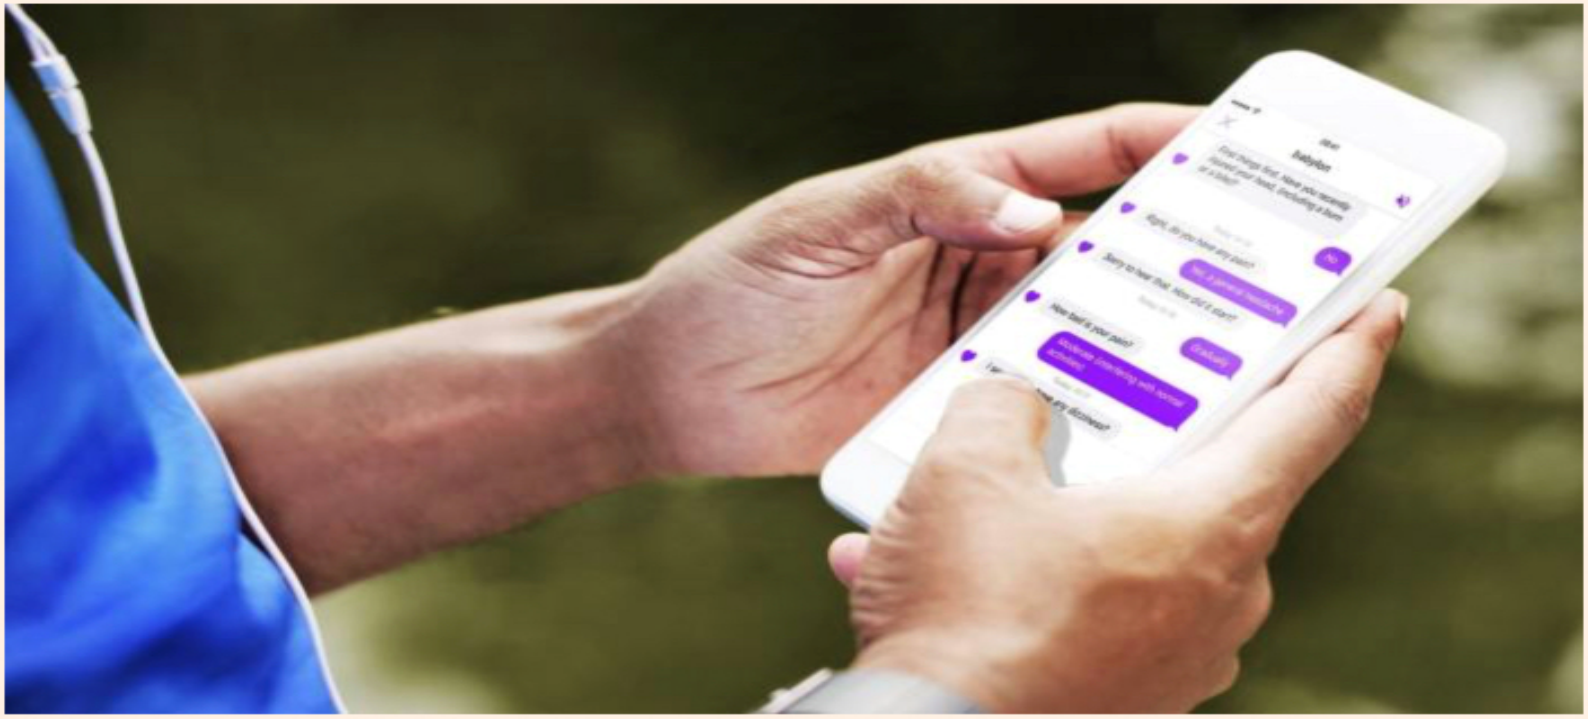

Babylon is one of a number of new tech products being adopted by overburdened health services

Aliya Ram and **Sarah Neville** JULY 13, 2018

41

Australian Alliance
for Artificial Intelligence in Healthcare

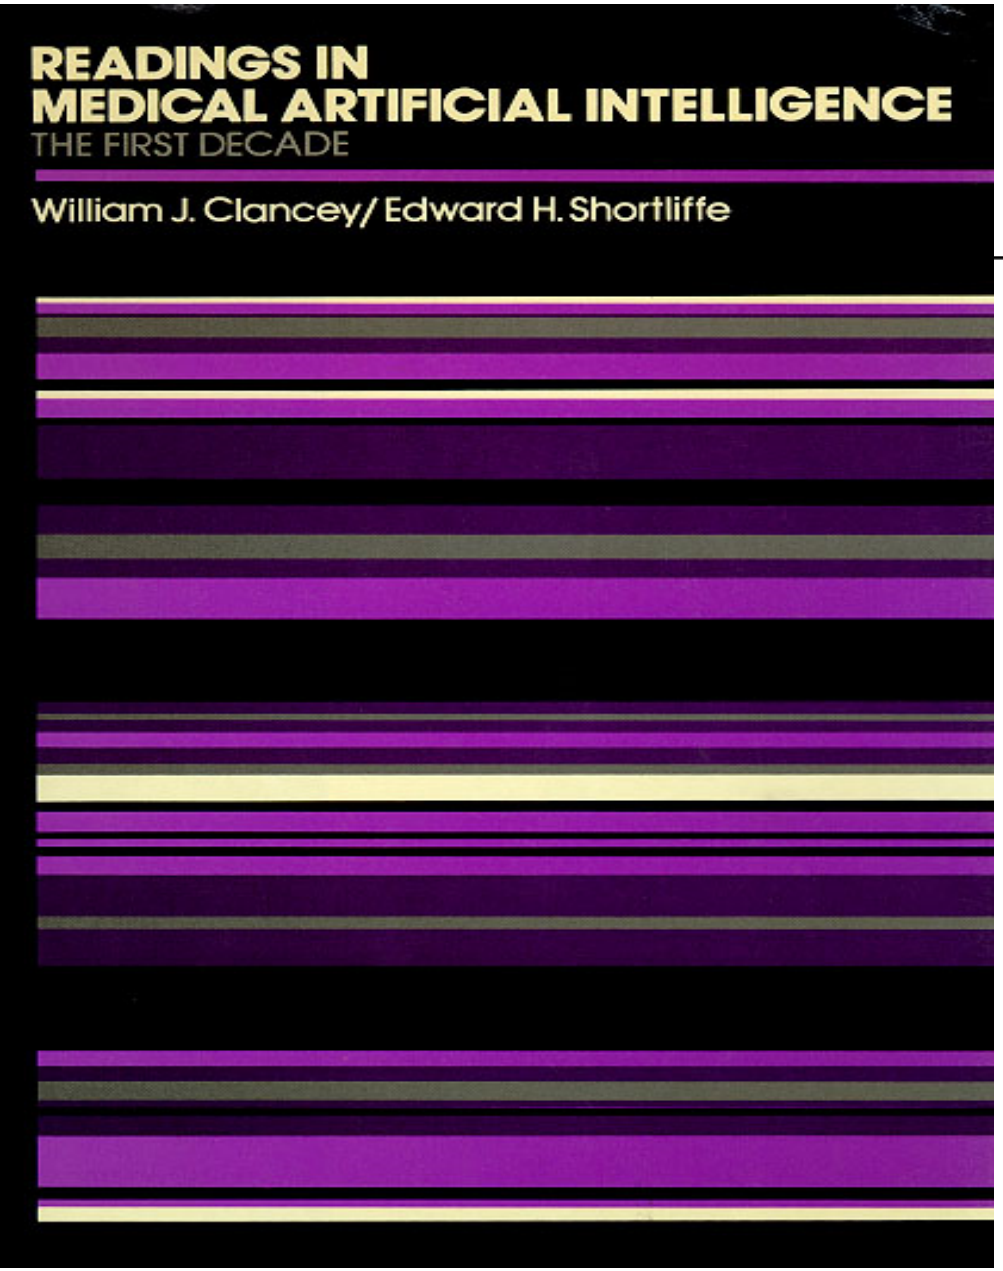

1943: McColloch and Pitts 1st paper on neural networks

1950: Turing test

1959: First AI in Medicine paper (Ledly and Lusted)

1963: First Radiology AI paper "Computer Diagnosis of Primary Bone Tumors" (*Radiology*, 80:2, 273)

1969: Minsky and Papert "kill" neural net research with their Perceptron book

1974: Lighthill Report kills AI research in the UK

1979: Symposium on AI in Medicine, American Association for Advancement of Science

1984: "First Decade" compilation text

1985: First AI in Medicine Conference (Pavia)

1985-90: Expert system "boom"

1991: First AI winter

2016-17: Deep Learning "boom"

2019-20: Is winter coming?

International Health AI Activity

Australian Alliance
for Artificial Intelligence in Healthcare

- English NHS allocated £1 billion to AI research/implementation, including goal of cutting cancer deaths by 10% (22,000 lives) per year by 2035
- Darzi review - automation could release GBP12.5 billion/10% of NHS budget
- ESPRC funding > 600 AI doctorates over 12 Centres of Excellence
- ~50% of AI healthcare start-ups are in UK

International AI Activity

Australian Alliance
for Artificial Intelligence in Healthcare

- China building a $US 2.1 billion AI technology park in Beijing's western suburbs. China’s AI industry forecast to be $150 billion in a decade.
- U.S. government spending on civilian AI programs in 2016 was about $US 1.2 billion
- MIT – new $1 billion endowment for transdisciplinary College of AI research
- Australia???? Minimal investment to date.

Australian Alliance
for Artificial Intelligence in Healthcare

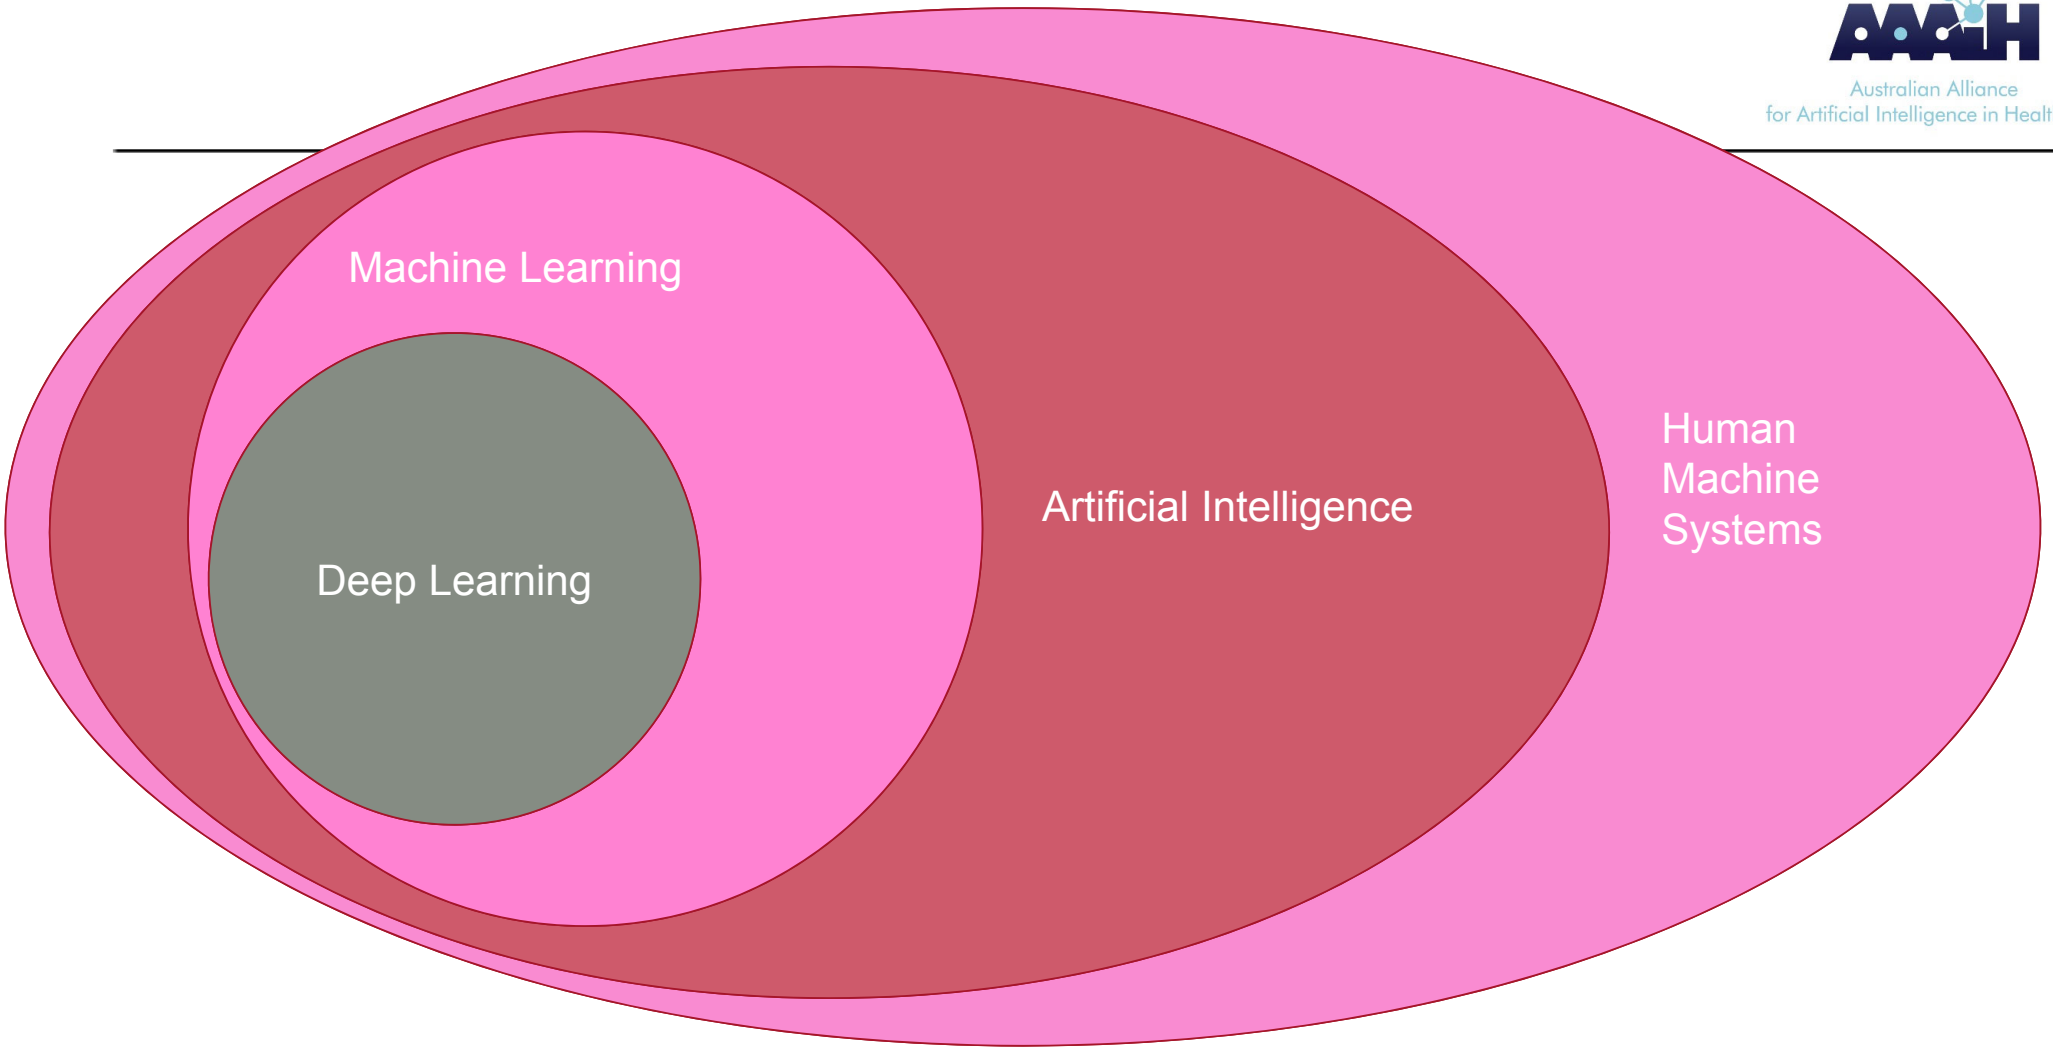

Different classes of AI

Australian Alliance for Artificial Intelligence in Healthcare

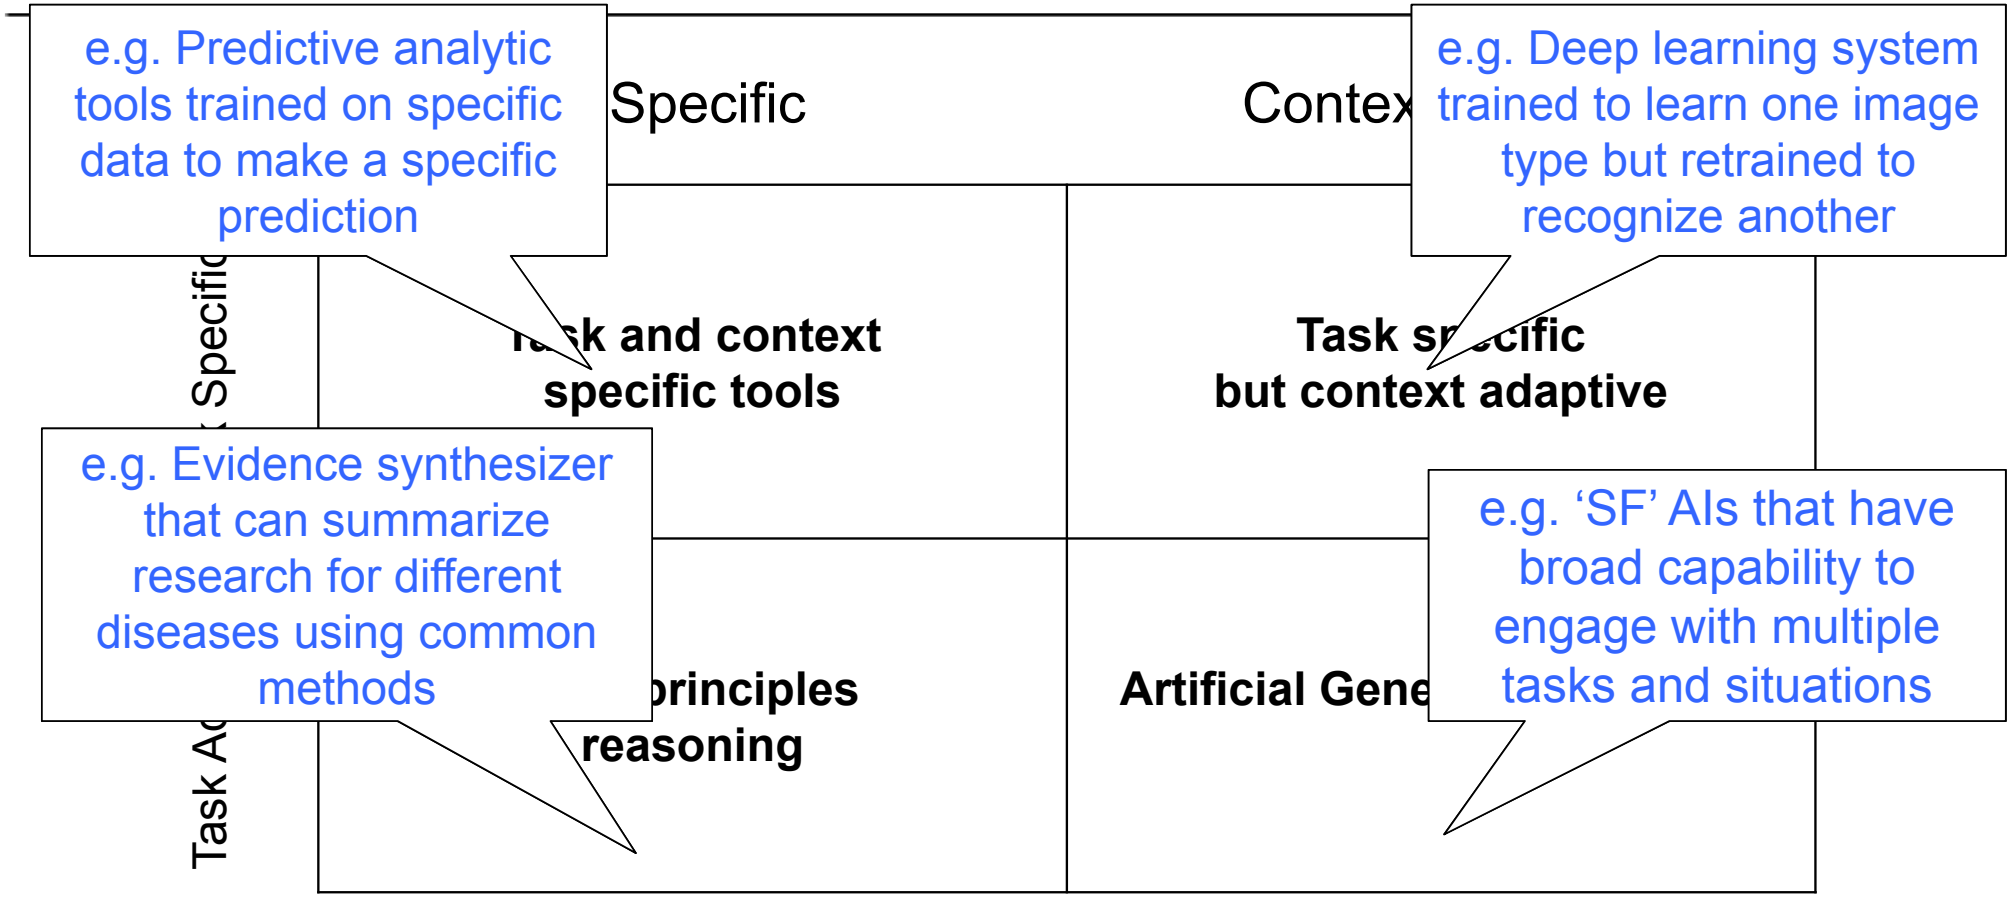

Different applications for AI in Healthcare

Australian Alliance for Artificial Intelligence in Healthcare

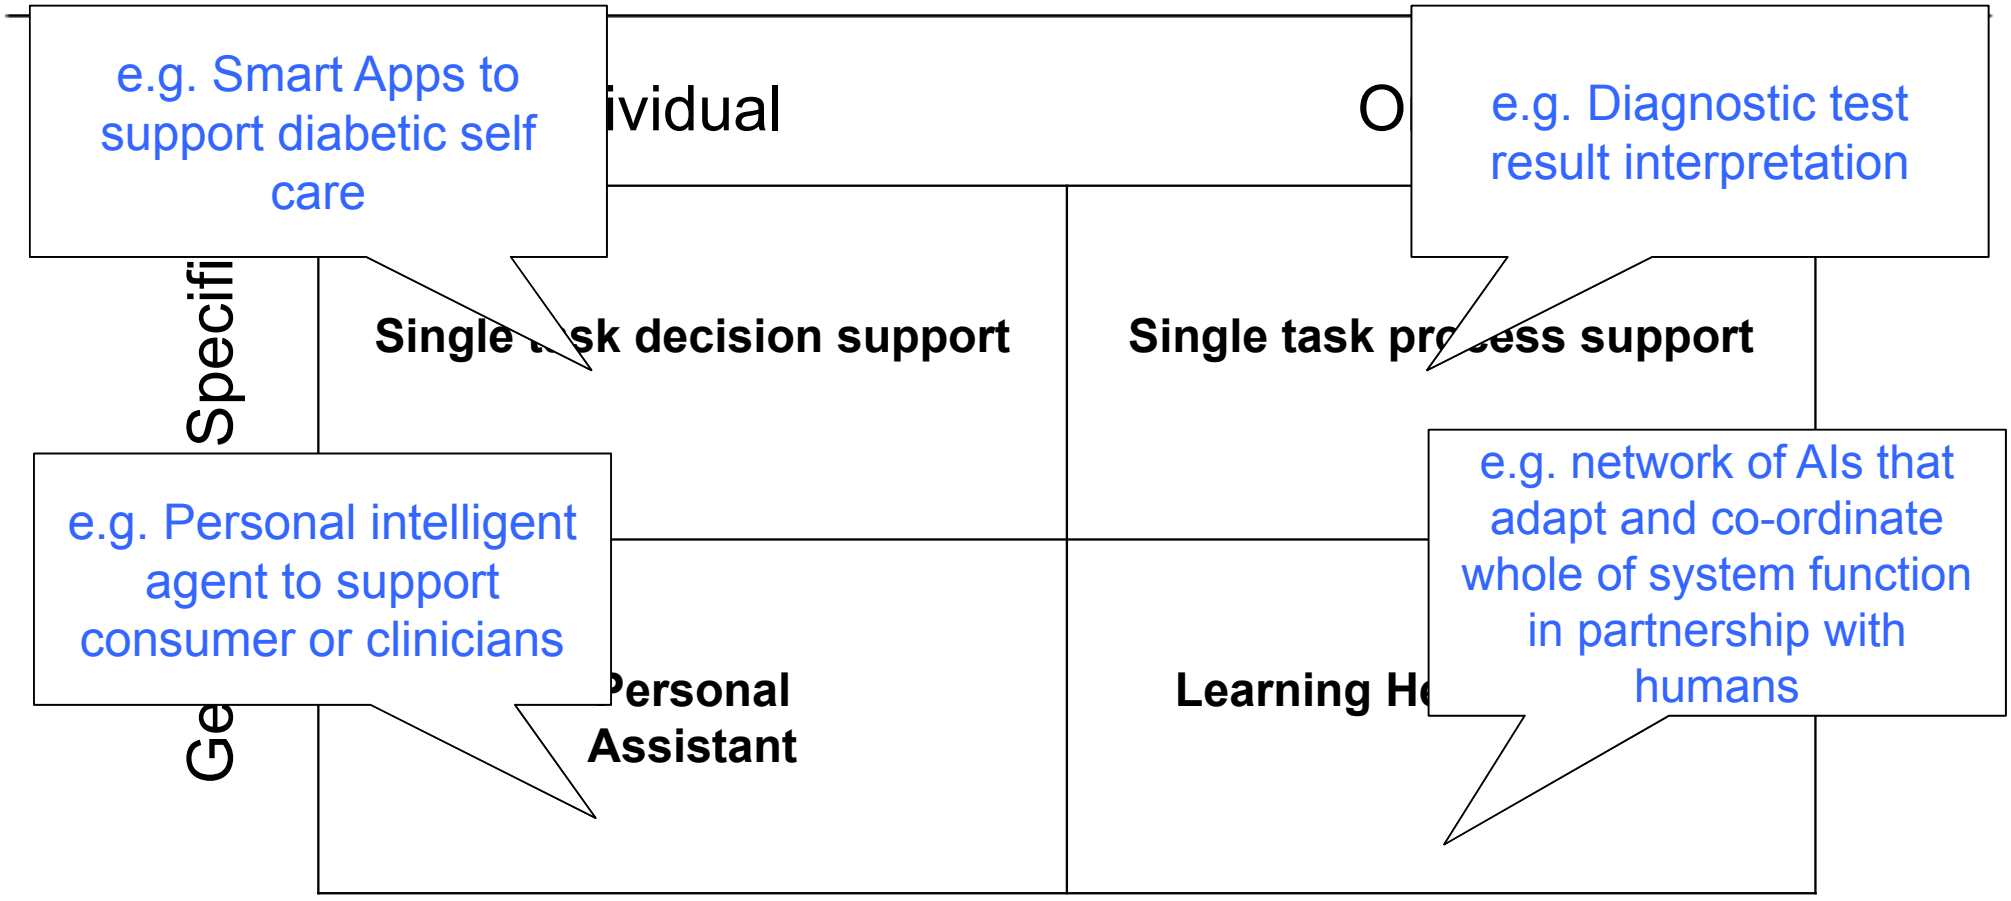

state of the art

Australian Alliance
for Artificial Intelligence in Healthcare

- Today's AIs are trained to do specific tasks with specific problems – diagnose specific illnesses, recommend specific therapies or investigations

- General-purpose artificial general intelligence (AGI) with across-task reasoning is not state of the art, and 10-20 years away at least

Object Detection, LSVRC Competition

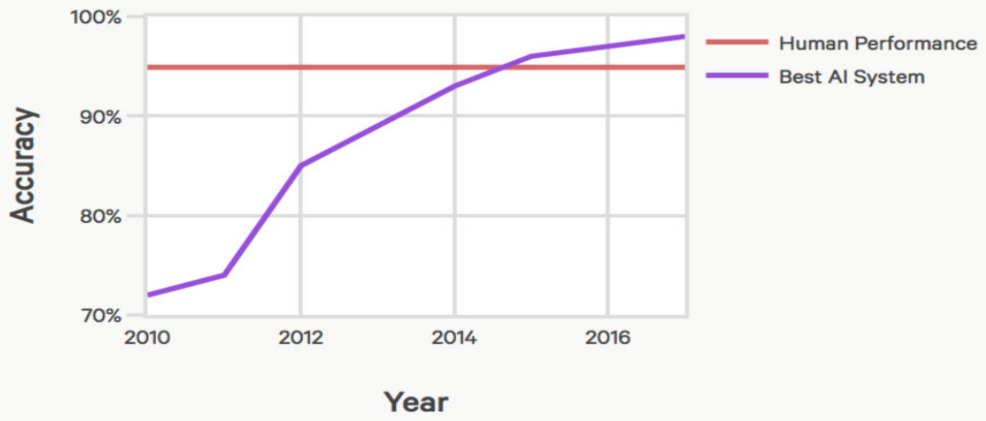

Speech Recognition, Switchboard HUB5'00

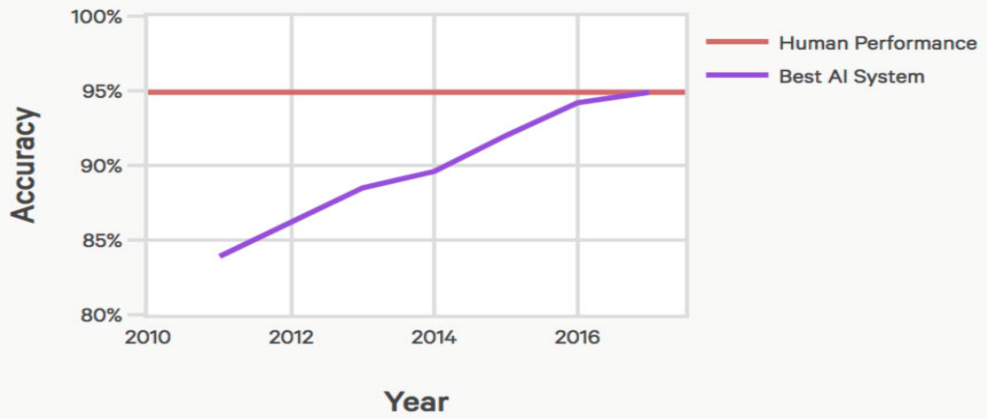

Question Answering, SQuAD v1.1

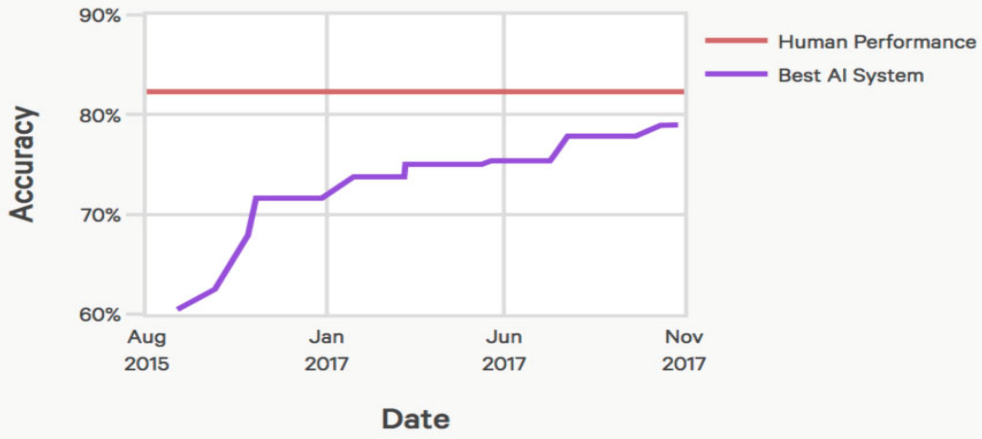

Source: stanford-qa.com

AIINDEX.ORG

Visual Question Answering, VQA 1.0

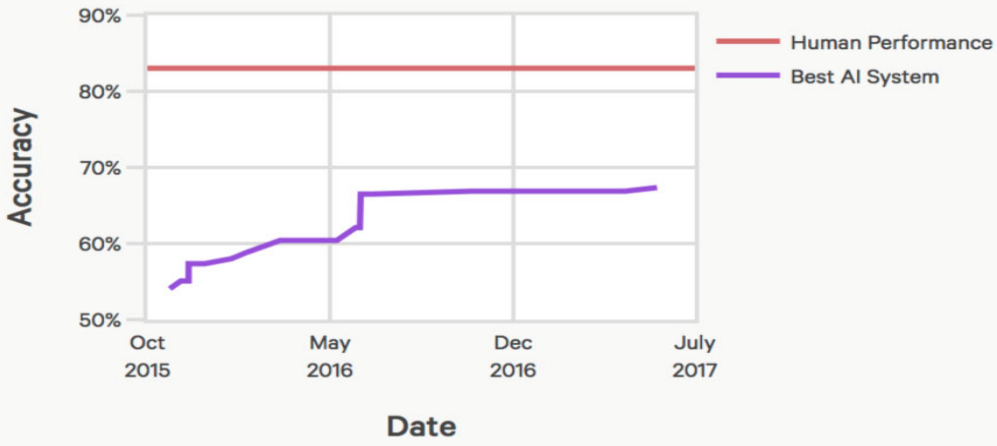

AIINDEX.ORG

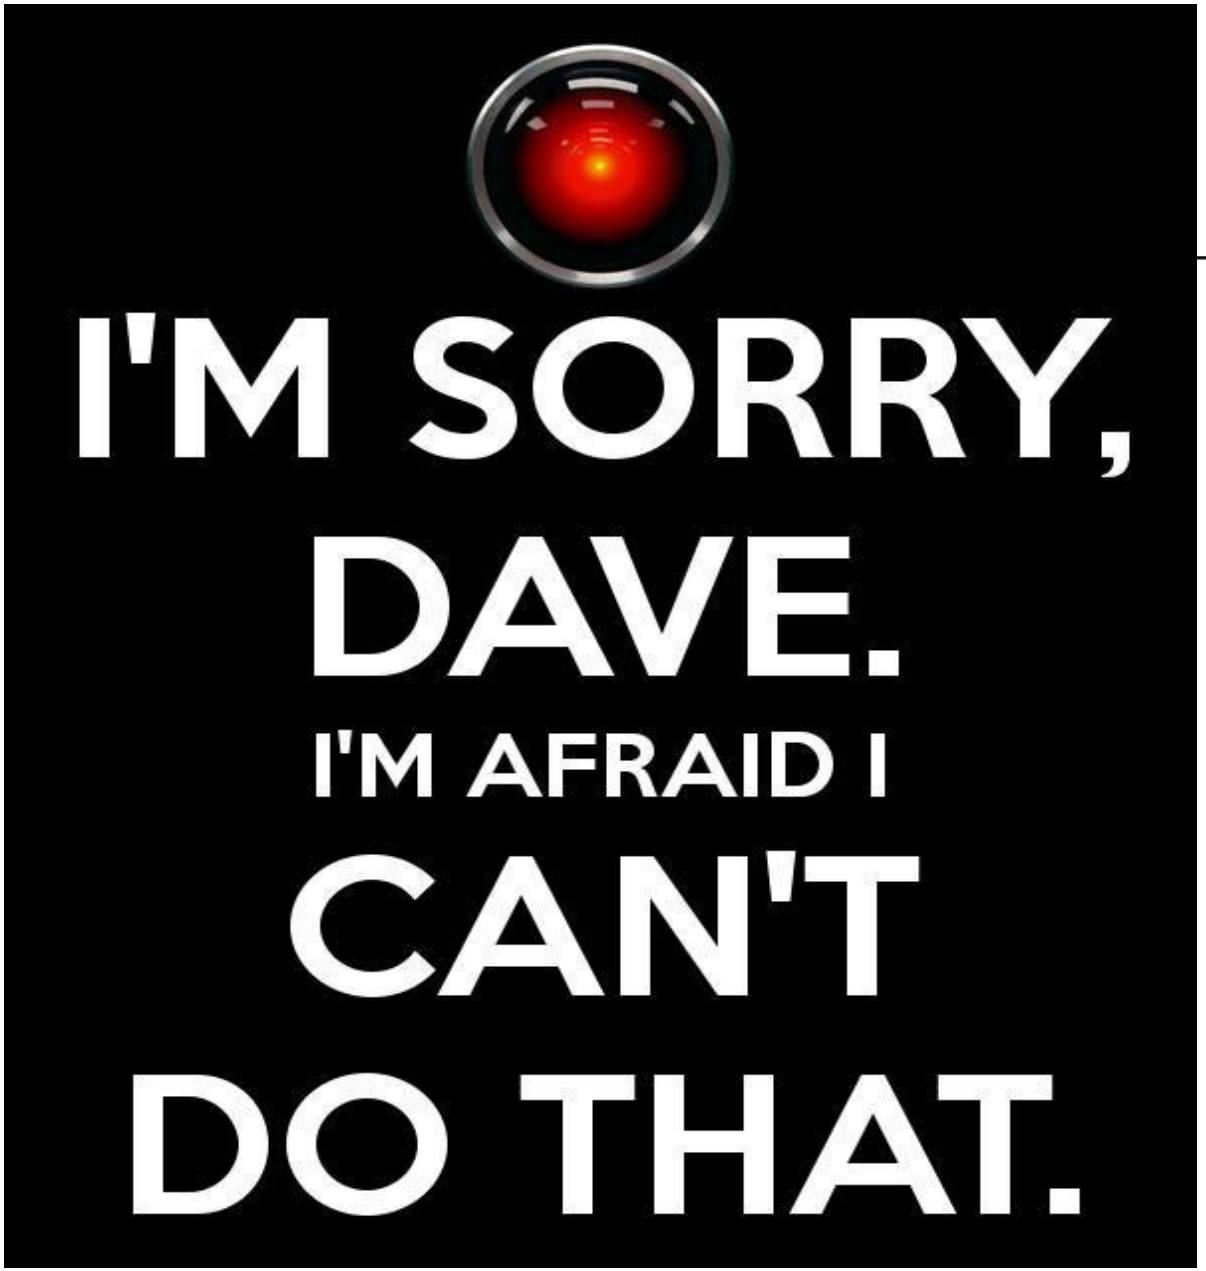

Australian Alliance
for Artificial Intelligence in Healthcare

Should we trust AI?

Asimov's Laws

Australian Alliance
for Artificial Intelligence in Healthcare

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws
4. A robot may not harm humanity, or, by inaction, allow humanity to come to harm.

Contents lists available at ScienceDirect

Resuscitation

journal homepage: www.elsevier.com/locate/resuscitation

Rapid response systems

Pre-existing risk factors for in-hospital death among older patients could be used to initiate end-of-life discussions rather than Rapid Response System calls: A case-control study☆

Magnolia Cardona-Morrell[a,*], Amanda Chapman[b], Robin M. Turner[c], Ebony Lewis[a], Blanca Gallego-Luxan[d], Michael Parr[a,b], Ken Hillman[a,b]

[a] *The Simpson Centre for Health Services Research, South Western Sydney Clinical School, The University of New South Wales, P.O. Box 6087 UNSW, Sydney, NSW 1466, Australia*

[b] *Intensive Care Unit, Liverpool Hospital, Level 2, Clinical Building, 1 Elizabeth St., Liverpool, NSW 2170, Australia*

[c] *School of Public Health and Community Medicine, The University of New South Wales, Level 2, Samuels Building, UNSW Kensington Campus, Sydney, NSW 2052, Australia*

[d] *Centre for Health Informatics, Australian Institute of Health Innovation, Macquarie University, Level 6, 57 Talavera Rd., North Ryde, NSW 2113, Australia*

Ethical challenges to AI in healthcare

Australian Alliance
for Artificial Intelligence in Healthcare

- Can an AI participate in end of life decisions without the risk of harm?
  - Is it ok to suggest withdrawal of care leading to death?
  - Is it ok to suggest continuation of care that is futile and leads to prolonged suffering?
- Can an AI triage patients to care or limited care? (e.g. Emergency Dept, battlefield, civil incident)
- How does an AI choose between harming individuals for the sake of the population, or vice versa? e.g. quarantine decisions

you and your AI together

Australian Alliance
for Artificial Intelligence in Healthcare

- Can be the best of both worlds:
  - Blend AI single task focus with human general purpose reasoning (but remember specific knowledge trumps reasoning)
  - Massively increase the number of events sensed
  - Circumvent human cognitive bottlenecks through AI mediated complexity reduction

Woman Drives for 900 Miles Instead of 90 Thanks to GPS Error

Jesus Diaz

1/14/13 12:14pm · Filed to: GPS

326.6K 483 2

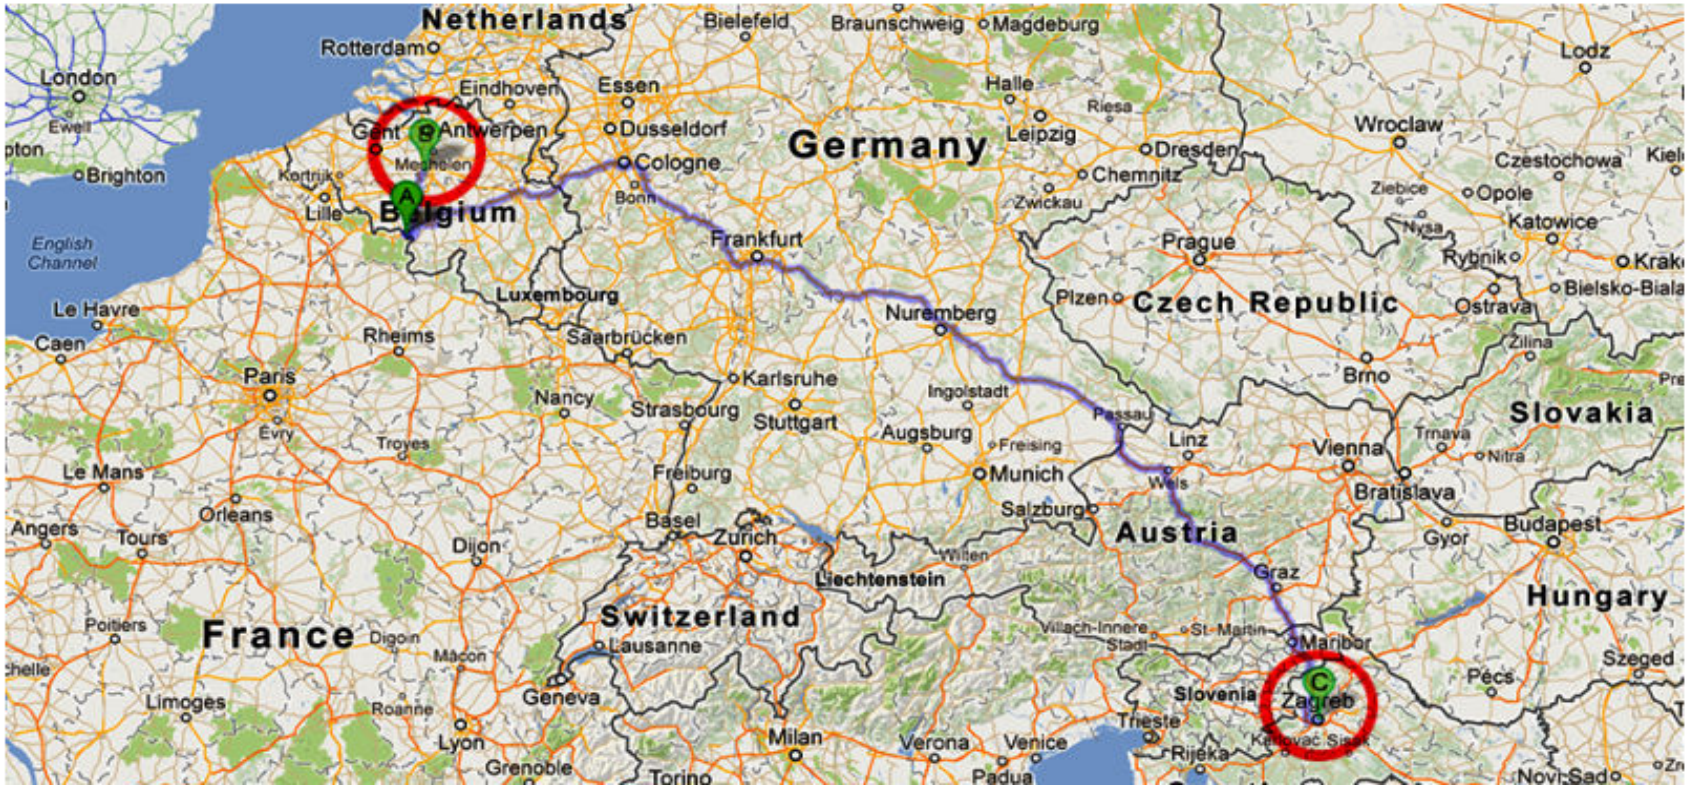

Man killed in Tesla auto-drive crash may have been watching Harry Potter

Barbara Liston and Bernie Woodall

A digital video disc player was found in the Tesla car that was on autopilot when its driver was killed in a May 7 collision with a truck, Florida Highway Patrol officials said on Friday.

Whether the portable DVD player was operating at the time of the crash has not been determined, however, and witnesses who came upon the wreckage of the 2015 Model S sedan gave differing accounts on Friday about whether the player was showing a movie.

Automation bias – Too much trust

Australian Alliance
for Artificial Intelligence in Healthcare

"The tendency to use automated cues as a heuristic replacement for vigilant information seeking and processing"

(Mosier & Skitka,1996)

- Arises when automation works well but not perfectly
- Also known as automation induced complacency
- Associated with reduced human vigilance

(Lyell and Coiera, JAMIA, 2016)

The Digital Health Trade-Off

Australian Alliance
for Artificial Intelligence in Healthcare

1: E-health can make patient care safer …

2: … but poorly designed, implemented, or used, E-health can lead to patient harm or death.

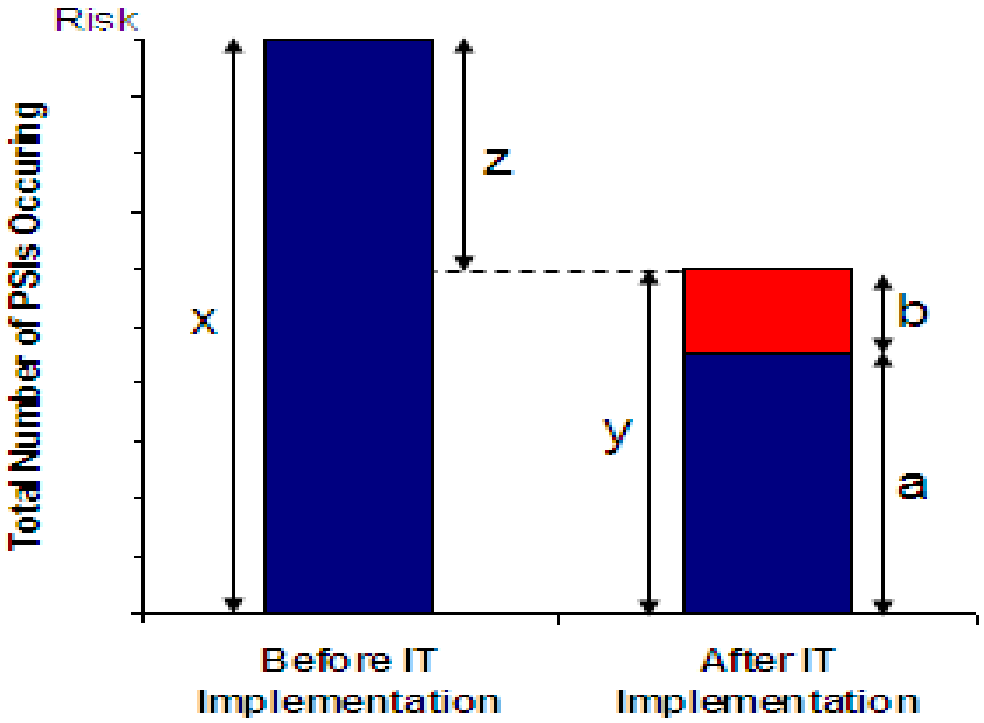

Safer care: x > y

x = no. of incidents before IT implementation

y = no. of incidents after IT implementation

a = incidents occurring despite, but not because of the new IT

b = incidents associated with new IT

Baker, M et al, Safer IT in a Safer NHS: Account of a Partnership, The British Healthcare Computing & Information Management, Vol. 23 No. 7 Sept 2006

work as imagined

work as done

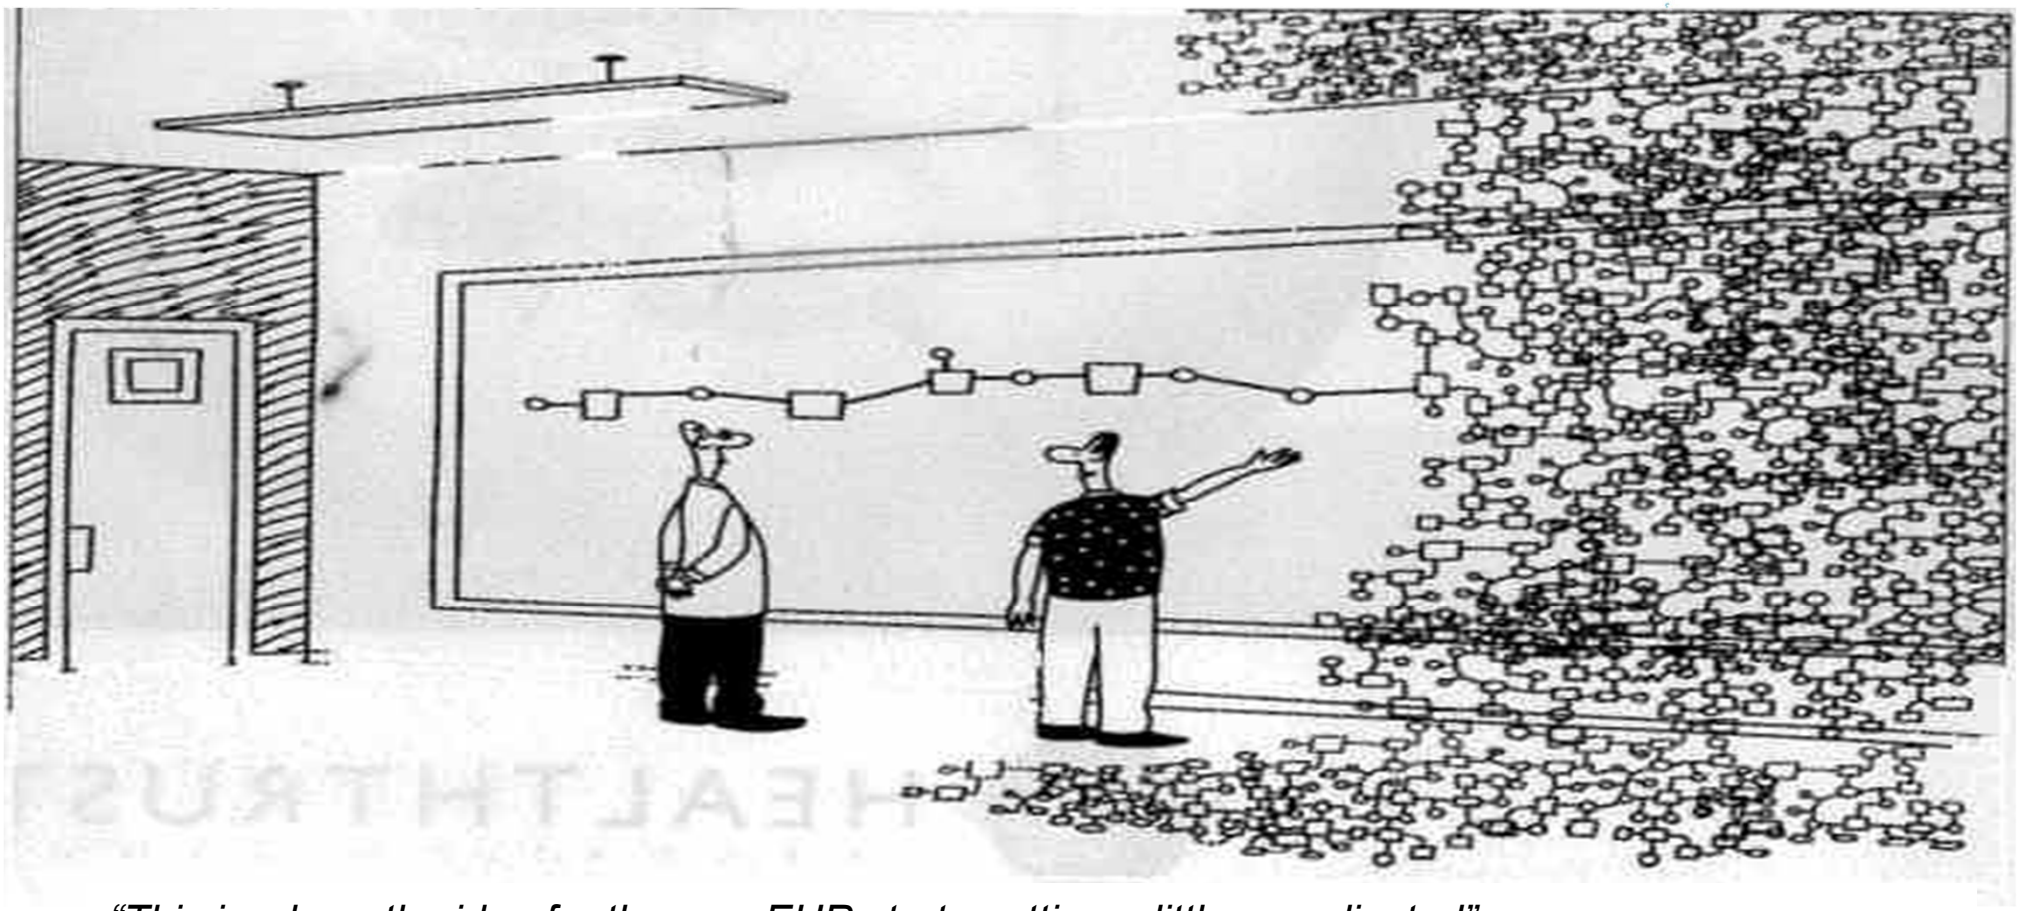

“This is where the idea for the new EHR starts getting a little complicated”

Implementation Shapes Outcome

Australian Alliance
for Artificial Intelligence in Healthcare

- Variations in digital health implementation outcome are the norm.
- Post-implementation analysis often reveals disappointing outcomes compared to expectations from lead sites (with different context).
- A 62 hospital CPOE study found capacity of the implemented systems to detect harmful medication orders varied widely even if using same system (Metzger, Welebob et al. 2010).
- Only 27% of variation associated choice of EHR.
- Results for using same CPOE varied by 40–65%.

The last mile

Australian Alliance
for Artificial Intelligence in Healthcare

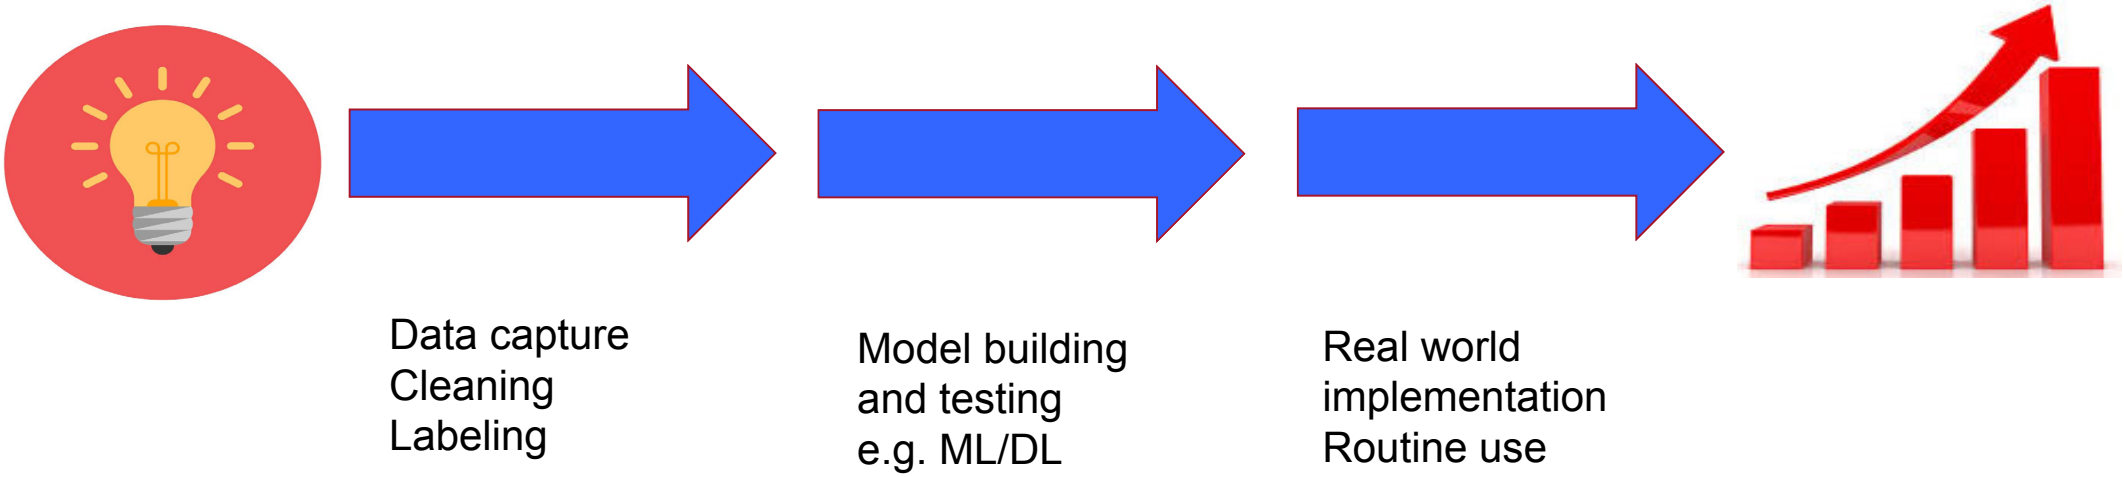

Beyond Deep Learning

- The **same algorithm** with the **same outputs** can have **substantially different real world effects** in different settings
- Last mile problems include bias in training data, implementation and context effects
- Working with AIs also requires solutions to problems in ethics, safety, human education, human-computer interaction design, and the behavior of human-computer systems.

Before you decide to trust a clinical AI recommendation …

Australian Alliance for Artificial Intelligence in Healthcare

- Ask yourself:
  - What do I think is going on? Form a pre-recommendation view. Otherwise you risk automation bias.
  - What is the benchmark performance for this AI on this disease? Is it good at this specific task?
  - Do you understand the reasoning behind the AI recommendation? If it is a black box then how do you know it is right?
  - Do you have access to data unavailable to the AI? If so, you will need to revise the recommendation in the light of this additional evidence.
  - Is the patient representative of the population on which the AI was trained? If not, then the AI recommendations are biased.
  - Is the patient presenting with complex multimorbidity? If so then there may be confounding that makes the recommendation misleading.
  - Are there differences in data collection methods at your site compared to those used to collect training data? If so, then the recommendation may be biased.

Amara's Law

Australian Alliance
for Artificial Intelligence in Healthcare

"We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run"

The fate of medicine in the time of AI

Australian Alliance
for Artificial Intelligence in Healthcare

- You wont see massive change in 5 yrs. You will in 10.
- Don’t expect human redundancy. Do expect transformation.
- Managing the safety vs. innovation debate critical.
- If AI gets a radiologist out of the reading room who will pay for patient-centered radiology?
- Will the diagnostic specialties merge? E.g. Radiology, Pathology.
- Disruption comes from left field e.g. AI assisted point of care imaging.

The Australian Alliance for Artificial Intelligence in Healthcare (AAAIH)

Australian Alliance
for Artificial Intelligence in Healthcare

AUSTRALIAN INSTITUTE OF HEALTH INNOVATION, MACQUARIE UNIVERSITY

An international partnership with national goals

Established	Commenced late 2017
Partners	52 formal partners to date
Government	5 State Departments of Health
Industry	Google Health, Accenture, Amazon AWS, IBM Research, Oracle, Orion, Alcidion
National	CSIRO, ADHA, Healthdirect Australia, ACSQHC, AIHW, Standards Australia
State	CEC, ACI, SA-NT Datalink
Academia	Universities: Macquarie, Melbourne, Queensland, Bond, UNSW, Sydney, South Australia, Adelaide, Flinders. MCRI, Australian Genomics, AIML
Professional	RANZCR, CHF, HISA, ACHI, MSIA
International	Stanford, Oxford, Toronto, Michigan, Liverpool

AI Alliance Goals

Australian Alliance
for Artificial Intelligence in Healthcare

- Support and accelerate the adoption of AI enabled health services in Australia
- Build a community of shared interest and pooled resource
- Undertake essential research needed to demonstrate AI-enabled services are scalable, sustainable, and have positive impacts on clinical and economic outcomes :
    - Precision Healthcare Flagship
    - Consumer Health Flagship
- Rapidly build the national workforce needed for success
- To ensure AI is safe and ethical, develop needed national regulatory and governance frameworks

Precision Healthcare Flagship

Australian Alliance
for Artificial Intelligence in Healthcare

- Making AI work in clinical care is a research and translational frontier
- Focus on projects that:
  - Tailor care to individuals
  - Exploiting data from EHRs, omics, wearables, social media ….
  - Using machine learning ..
  - Computational reasoning
  - Conversational agents
- Emphasis on implementation science and translational research
- Exemplars: Real time prediction of patient deterioration (Qld); Predicting Glaucoma progression (SA); Personalised discharge support (Vic); Cystic Fibrosis personalisation of care (NSW)

Consumer Health Flagship

Australian Alliance
for Artificial Intelligence in Healthcare

- Creating AI driven personal health navigator services for Australians
- Assist where there is system complexity, support self care, maximize prevention
- Exploit personalized care technologies
- Full consumer engagement in the co-design process
- Exemplars: recommend aged care services (Healthdirect Australia); AI for personalized behavior change for CVD (NSW); Conversational agents for drug free self management of chronic pain (SA); Safety of symptom checker (NSW)

Workforce program

Australian Alliance
for Artificial Intelligence in Healthcare

- Target groups:
  - Researchers
  - Health service leaders
  - Clinicians

- Through its membership, facilitate and develop educational programs:
  - AI doctoral and masters programs
  - AI components within existing professional programs e.g. via Professional Colleges
  - Pre-qualification training e.g. training for safe and effective use of AI and digital health in medical curriculum

AI Alliance work plan 2019

Australian Alliance
for Artificial Intelligence in Healthcare

✓ State hubs activated in Queensland, NSW, Victoria, and SA – data, ethics, trial clinical sites, consumer and clinician engagement, feasibility assessments

✓ Working groups formed in:

- AI safety, quality and ethics
- AI enabled workforce
- Precision Healthcare Flagship
- Consumer Health Flagship

✓ Secure funding for first round of Flagship Programs from 2020

✓ 2020-2024 work program agreed and finalized

What is the fate of medicine in the time of AI?

Our fate is to change.

(*The Lancet*, 392, 2331, 2018)

AIHI | AUSTRALIAN INSTITUTE OF HEALTH INNOVATION

MACQUARIE University

Thank you

@enricocoiera

enrico.coiera@mq.edu.au