Thinking about Thinking with AI

BE Chapman, PhD

2023-09-28

Acknowledgements

- Simon D'Alfonso
- Mike Conway
- Bruce Buchanan
- Students at University of Pittsburgh
- Students at UCSD
- Students at University of Melbourne

AI in Healthcare

- Initial attitudes towards AI in healthcare
- `PollEv.com/brianchapman270`
- `brianchapman270 to 22333`

AI in Healthcare

- “The most profoundly depressing of all ideas about the future of the human species is the concept of artificial intelligence.” (Lewis Thomas, M.D.)
  - New England Journal of Medicine
  - February 28, 1980 (“Artificial Intelligence” 1980)

Section 1

What do we mean by Artificial Intelligence?

What I was taught!

"Artificial intelligence is mostly artificial and not very intelligent."

Figure 1: T. Allan Pryor

What is AI?

The field of artificial intelligence (AI) officially started in 1956, launched by a small but now-famous DARPA-sponsored summer conference at Dartmouth College, in Hanover, New Hampshire. . . .From where we stand now, into the start of the new millennium, the Dartmouth conference is memorable for many reasons, including. . . the term 'artificial intelligence' was coined there. . . .(Bringsjord and Govindarajulu 2022)

What is AI?

Though the term 'artificial intelligence' made its advent at the 1956 conference, certainly the field of AI, operationally defined... was in operation before 1956. (Bringsjord and Govindarajulu 2022)

Precursors of AI: Gottfried Wilhelm Leibniz

When controversies arise, there will be no more need for a disputation between two philosophers than there would be between two accountants [computistas]. It would be enough for them to pick up their pens and sit at their abacuses, and say to each other (perhaps having summoned a mutual friend): 'Let us calculate.' (quoted in (Bringsjord and Govindarajulu 2022))

[a]

[a] Wikipedia

What is in a name? Herbert Simon

What if instead of "artificial intelligence" we said "complex information processing"?

"What's in a name? That which we call a rose, by any other word would smell as sweet." ("Romeo and Juliet")

[a]

[a] Wikipedia

The term “Artificial Intelligence” is marketing

What is AI?[1]

	Human-based	Ideal Rationality
Reasoning-Based	Systems that think like humans.	Systems that think rationally.
Behavior-Based	Systems that act like humans.	Systems that act rationally.

[1](Bringsjord and Govindarajulu 2022)

AI or not AI?

- AI or not AI?
- `PollEv.com/brianchapman270`
- `brianchapman270 to 22333`

AI or not AI?

John McCarthy: “As soon as it works, no one calls it AI anymore.” (Mitchell 2019)

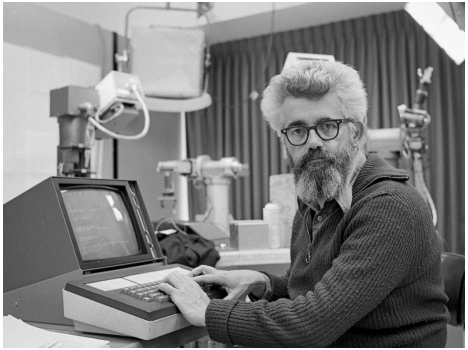

[a]

[a] The Independent

What is “Artificial”?

- Created by humans through engineering processes
  - Training pigeons to recognize cancerous cells in histology slides (Fry 2019, 79) would not be artificial
  - Nor would training dogs to recognize GI cancers through smell
  - But building cell-based computers might be

What is “Intelligence”?

- Multifaceted and complex cognitive ability that encompasses various mental faculties.
- Capacity to acquire, process, and apply knowledge to:
  - reason, solve problems
  - abstract, conceptualise, comprehend
  - make decisions
  - adapt to new situations
  - learn from experiences

What is “Intelligence”?

- Intelligence can manifest in various forms:
  - logical-mathematical intelligence
  - linguistic intelligence
  - spatial intelligence
  - emotional intelligence

Three flavors of AI (in alphabetical order)

- Logic/Expert Systems
  - Represent knowledge as computable rules and relationships
- Machine Learning
  - Learn implicit or explicit relationships between data (input) and classifications, predictions, etc. (output)
- Probability
  - Represent knowledge as as probabilities and use Bayes' theorem to make inferences

Section 2

Why AI?

What are brains good for?

What, then, is the role of the biological brain. . . .It is expert at recognizing patterns, at perception, and at controlling physical actions, but it is not so well designed. . . for complex planning and long, intricate, derivations of consequences. ***It is, to put it bluntly, bad at logic and good at Frisbee.*** *(Clark 2003)*

Human brain exceptionalism

For what is special about human brains, and what best explains the distinctive features of human intelligence, is precisely their ability to enter into deep and complex relationships with nonbiological constructs, props, and aids. This ability, however, does not depend on physical wire-and-implant mergers, so much as on our openness to information-processing mergers. (Clark 2003)

“Information-processing mergers” (Clark 2003)

- External objects that aid the mind’s reasoning
  - Example: paper and pencil to do arithmetic
  - Example: Cell phones
    - “Finnish youngsters have dubbed the cell phone ‘kanny,’ which means extension of the hand.”

Natural-born Cyborgs

- Our *minds* think with
- Our biological brains

Natural-born Cyborgs

- Our *minds* think with
- Our biological brains
- Our bodies

Natural-born Cyborgs

- Our *minds* think with
- Our biological brains
- Our bodies
- External objects

Natural-born Cyborgs

- Our *minds* think with
- Our biological brains
- Our bodies
- External objects

Extended mind theory

Medicine makes extensive use of extended senses

- Diagnostic imaging
- Stethescopes
- Microscopes

Poll Everywhere

- Is the extended mind theory different?
- `PollEv.com/brianchapman270`
- `brianchapman270 to 22333`

Poll Everywhere

- Should we think about what we are doing?
- `PollEv.com/brianchapman270`
- `brianchapman270 to 22333`

Limited Cognitive Capacity: Alfred North Whitehead

It is a profoundly erroneous truism, repeated by all copy-books and by eminent people when they are making speeches, that we should cultivate the habit of thinking of what we are doing. The precise opposite is the case. ***Civilization advances by extending the number of important operations which we can perform without thinking about them.*** *Operations of thought are like cavalry charges in a battle — they are strictly limited in number, they require fresh horses, and must only be made at decisive moments. (Whitehead 1958)*

Section 3

Why AI in healthcare?

Francis Bacon: Idols of the Mind [2]

- Universal mental limitations “inherent in human nature”
- Acquired beliefs of human problem solvers “corrupts the light of nature”
- The limits of language “throw everything into confusion”
- Developed systems of thought “which have become inveterate by tradition, implicit credence and neglect.”[a]

[a] “For the past 40 years, research by psychologists has confirmed the writings of Bacon.” Weed

[a] Wikipedia

[2] This summarization of Bacon is from Larry Weed.

Why did early AI researchers work on medical problems?

- Because it is interesting
- Because the practice of modern medicine exceeds the capacity of the unaided human mind

Healthcare is expensive [3]

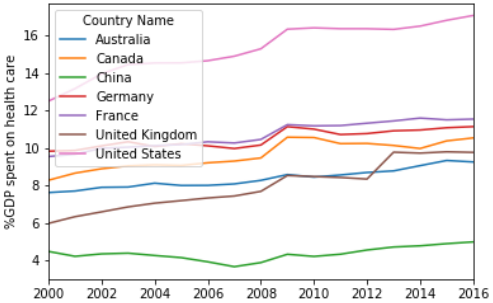

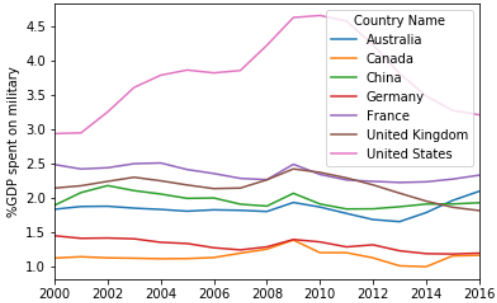

[3] Data are taken from the World Bank

Healthcare is dangerous

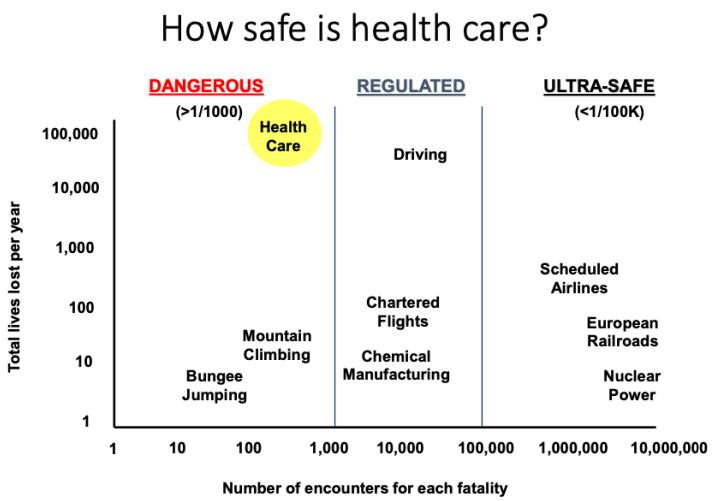

Harvard School of Public Health

Healthcare has delivery challenges (Braithwaite, Glasziou, and Westbrook 2020)

Braithwaite *et al. BMC Medicine* (2020) 18:102
https://doi.org/10.1186/s12916-020-01563-4

Received: 30 July 2019 Revised: 11 March 2020
Accepted: 17 March 2020 Published online: 04 May 2020

The three numbers you need to know about healthcare: the 60-30-10 Challenge

Jeffrey Braithwaite[1*], Paul Glasziou[2] and Johanna Westbrook[3]

Abstract

Background: Healthcare represents a paradox. While change is everywhere, performance has flatlined. 60% of care on average is in line with evidence- or consensus-based guidelines, 30% is some form of waste or of low value, and 10% is harm. The 60-30-10 Challenge has persisted for three decades.

Main body: Current top-down or chain-logic strategies to address this problem, based essentially on linear models of change and relying on policies, hierarchies, and standardisation, have proven insufficient. Instead, we need to marry ideas drawn from complexity science and continuous improvement with proposals for creating a deep learning health system. This dynamic learning model has the potential to assemble relevant information including patients' histories, and clinical, patient, laboratory, and cost data for improved decision-making in real time, or close to real time. If we get it right, the learning health system will contribute to care being more evidence-based and less wasteful and harmful. It will need a purpose-designed digital backbone and infrastructure, apply artificial intelligence to support diagnosis and treatment options, harness genomic and other new data types, and create informed discussions of options between patients, families, and clinicians. While there will be many variants of the model, learning health systems will need to spread, and be encouraged to do so, principally through diffusion of innovation models and local adaptations.

Conclusion: Deep learning systems can enable us to better exploit expanding health datasets including traditional and newer forms of big and smaller-scale data, e.g. genomics and cost information, and incorporate patient preferences into decision-making. As we envisage it, a deep learning system will support healthcare's desire to continually improve, and make gains on the 60-30-10 dimensions. All modern health systems are awash with data, but it is only recently that we have been able to bring this together, operationalised, and turned into useful information by which to make more intelligent, timely decisions than in the past.

Keywords: Learning health system, Complexity, Complexity science, Change, Evidence-based care, Clinical networks, Quality of care, Patient safety, Policy, Healthcare systems

Healthcare delivery challenge [4]

Does AI in healthcare need to look the same everywhere?

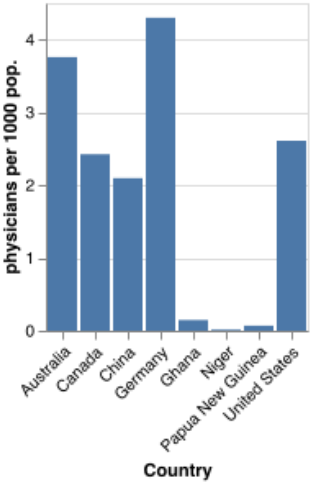

[4] Data are taken from the World Bank

Overwhelming amount of information (Medicine et al. 2008)

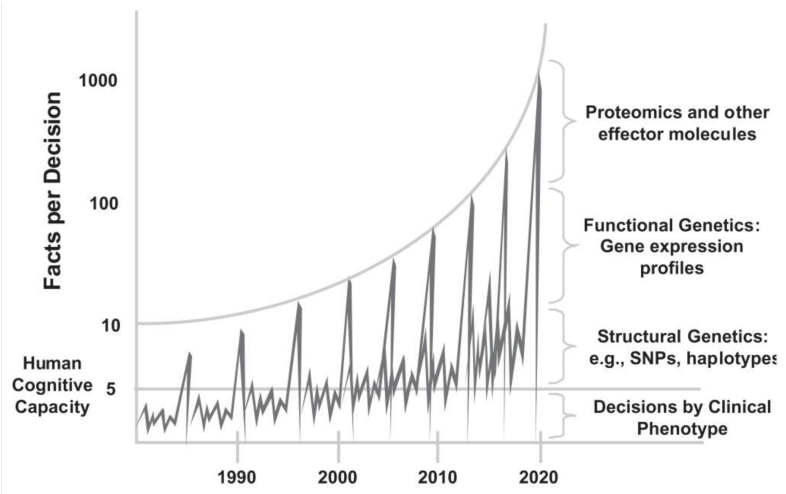

Early Probabilistic AI in Healthcare: 1961

"A Mathematical Approach to Medical Diagnosis: Application to Congenital Heart Disease" (WARNER et al. 1961)

Early Expert System AI in Healthcare: MYCIN (1973)

"An artificial intelligence program to advise physicians regarding antimicrobial therapy" (Shortliffe et al. 1973)

Early Expert System AI in Healthcare: MYCIN

But with MYCIN, the program that diagnoses disease and prescribes treatment, the experts were not willing to do away with the doctor and let the computer dispense pills to each patient in the hospital.

"If a program such as MYCIN were acting independently of a physician, I think that would be inappropriate," said Professor Buchanan, **who has a computer terminal in his home.** "The problem is that human diseases are open-ended.

Experts Argue Whether Computers Could Reason, and if They Should

By LEE DEMBART

Can machines think? Should they? The computer world is in the midst of a fundamental dispute over those questions because an eminent computer scientist wrote a book arguing that machines could never be made to reason like people and, what was more, should not be.

Twenty years ago, in the infancy of the computer revolution, before problems began cropping up, the public was told that computers would be smarter than brains. Computer chess champions and machine translation, for example, were just around the corner. So far, neither has been accomplished successfully, and neither is likely to be any time soon.

Nor have computers had much success in making decisions that require judgment. They can rattle off the Manhattan telephone directory unerringly time after time, which no human can do, but they cannot begin to distinguish one face from another, as babies can.

Computer scientists have always said, "Give us more time. The problem is more complex than we thought." Then one of them, Joseph Weizenbaum, a professor of computer science at the Massachusetts Institute of Technology, wrote a book saying that the project was fundamentally unsound and dangerous to pursue, partly, he said, because the computers' and humans' ways of thought would always be alien, and because knowledge might become limited to what a computer could undertand.

The elders of the artificial intelligence community reacted with outrage. Even those who agreed with his premises criticized the book as being too harsh in tone, too personal in its attacks. Computer journals have bristled over the last year with reviews, comments and replies provoked by Professor Weizenbaum's book, "Computer Power and Human Reason" (W. H. Freeman & Co., 1976). Now the controversy has spilled into the prestigious publication Science,

Continued on Page 34, Column 3

erals of his intention to back Mr. Cuomo all the way in an apparent effort to strengthen Mr. Cuomo's bid for the Liberal Party's nomination.

Titular Head Turns Elsewhere

If Mr. Cuomo wins the Liberal nomination but loses that of his own party, it would mean that the Governor would be supporting the Liberal candidate for Mayor rather than the Democratic one.

When Mr. Boone was asked about the Carey move, disclosed in The New York Times yesterday, he said it was difficult to believe "the titular head of the Democratic Party would support a candidate of another party rather than his own, especially after all his pretensions about the importance of an open primary."

"Apparently the Liberal Party is more important to him than the Democratic

Continued on Page 23, Column 1

Casinos Bring Atlantic City Woes

By MARTIN WALDRON

Special to The New York Times

ATLANTIC CITY, May 7—Fires, the condemnation of slum housing and the destruction of small hotels are driving out many elderly and poor as this seaside resort begins to experience sociological consequences in a prelude to the opening of gambling casinos.

Under a measure in the New Jersey Legislature, casines could be opened only in large hotels. State and city officials hope to rebuild the deteriorating city, once the nation's premier convention center, by attracting $500 million worth of these new facilities.

Hotels and other businesses may be erected in an area now by slum housing, much of sought by speculators.

City officials deny tha policy to try to drive the poor from Atlan tic City. A third of the city's 40,000 per manent residents are on welfare, wil unemployment over 20 percent.

About 10,000 of these—mostly black and Hispanic residents—live in the are expected to be in demand soon for th construction of casinos, restaurants an parking facilities.

In all, some 700 to 1,000 people hav

Continued on Page 36, Column 3

Front P

SHOW SU

Figure 2: NYT 1977

Early Expert System AI in Healthcare: MYCIN

"Our goal is to build a program that can assist working scientists with reasoning problems. You wouldn't expect a tool to have all of the power of a working scientist, **but you would hope that you could design a smart system to provide some of the solutions to subproblems.**" (D. 1977)

Experts Argue Whether Computers Could Reason, and if They Should

By LEE DEMBART

Can machines think? Should they? The computer world is in the midst of a fundamental dispute over those questions because an eminent computer scientist wrote a book arguing that machines could never be made to reason like people and, what was more, should not be.

Twenty years ago, in the infancy of the computer revolution, before problems began cropping up, the public was told that computers would be smarter than brains. Computer chess champions and machine translation, for example, were just around the corner. So far, neither has been accomplished successfully, and neither is likely to be any time soon.

Nor have computers had much success in making decisions that require judgment. They can rattle off the Manhattan telephone directory unerringly time after time, which no human can do, but they cannot begin to distinguish one face from another, as babies can.

Computer scientists have always said, "Give us more time. The problem is more complex than we thought." Then one of them, Joseph Weizenbaum, a professor of computer science at the Massachusetts Institute of Technology, wrote a book saying that the project was fundamentally unsound and dangerous to pursue, partly, he said, because the computers' and humans' ways of thought would always be alien, and because knowledge might become limited to what a computer could understand.

The elders of the artificial intelligence community reacted with outrage. Even those who agreed with his premises criticized the book as being too harsh in tone, too personal in its attacks. Computer journals have bristled over the last year with reviews, comments and replies provoked by Professor Weizenbaum's book, "Computer Power and Human Reason" (W. H. Freeman & Co., 1976). Now the controversy has spilled into the prestigious publication Science,

Continued on Page 34, Column 3

erals of his intention to back Mr. Cuomo all the way in an apparent effort to strengthen Mr. Cuomo's bid for the Liberal Party's nomination.

Titular Head Turns Elsewhere

If Mr. Cuomo wins the Liberal nomination but loses that of his own party, it would mean that the Governor would be supporting the Liberal candidate for Mayor rather than the Democratic one.

When Mr. Buono was asked about the Carey move, disclosed in The New York Times yesterday, he said it was difficult to believe "the titular head of the Democratic Party would support a candidate of another party rather than his own, especially after all his pretensions about the importance of an open primary."

"Apparently the Liberal Party is more important to him than the Democratic

Continued on Page 23, Column 1

Casinos Bring Atlantic City Woes

By MARTIN WALDRON

Special to The New York Times

ATLANTIC CITY, May 7—Fires, the condemnation of slum housing and the destruction of small hotels are driving out many elderly and poor as this seaside resort begins to experience sociological consequences in a prelude to the opening of gambling casinos.

Under a measure in the New Jersey Legislature, casines could be opened only in large hotels. State and city officials hope to rebuild the deteriorating city, once the nation's premier convention center, by attracting $500 million worth of these new facilities.

Hotels and other businesses may be erected in an area now by slum housing, much of sought by speculators.

City officials deny tha policy to try to drive the poor from Atlan tic City. A third of the city's 40,000 per manent residents are on welfare, wit unemployment over 20 percent.

About 10,000 of these—mostly black and Hispanic residents—live in the are expected to be in demand soon for th construction of casinos, restaurants an parking facilities.

In all, some 700 to 1,000 people hav

Continued on Page 36, Column 3

Front P

SHOW SL

Figure 3: NYT 1977

Early Machine Learning AI in Healthcare (1985)

"A program for machine learning of counting criteria: empirical induction of logic-based classification rules" (Spackman 1985)

Section 4

Challenges with AI

Challenges

- Malicious actors
  - Adversarial attacks
- Imperfect models

Malicious Actors (Mirsky et al. 2019)

CT-GAN: Malicious Tampering of 3D Medical Imagery using Deep Learning

Yisroel Mirsky[1], Tom Mahler[1], Ilan Shelef[2], and Yuval Elovici[1]
[1]Department of Information Systems Engineering, Ben-Gurion University, Israel
[2]Soroka University Medical Center, Beer-Sheva, Israel
yisroel@post.bgu.ac.il, mahlert@post.bgu.ac.il, shelef@bgu.ac.il, and elovici@bgu.ac.il

Published in the 28th USENIX Security Symposium (USENIX Security 2019)

Demo video with pen-test: https://youtu.be/_mkRAArj-x0
Source code and datasets: https://github.com/ymirsky/CT-GAN

Malicious Actors

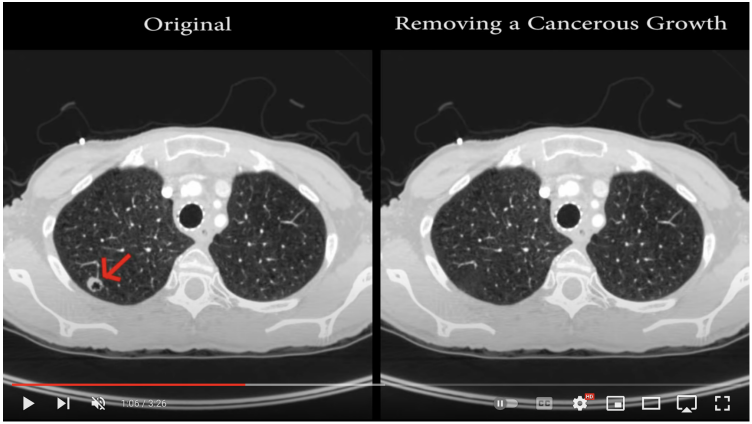

Injecting and Removing Cancer from CT Scans

Cyber Security Labs @ Ben Gurion University
6.74K subscribers
Subscribe

160 Share Clip Save

Adversarial Example (Mitchell 2019)

FIGURE 18: Original and "adversarial" examples for AlexNet. The left image

Adversarial Example (Mitchell 2019)

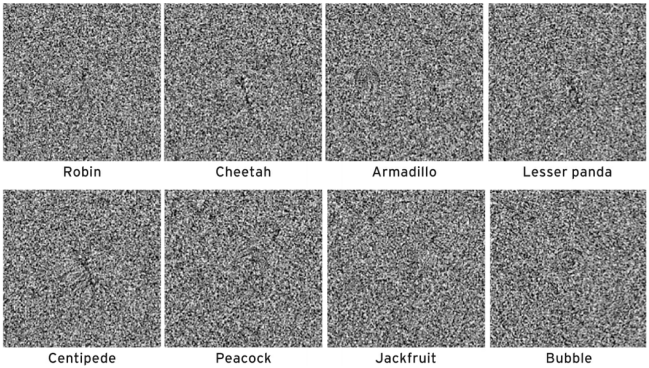

FIGURE 19: Examples of images created by a genetic algorithm specifically to fool a convolutional neural network. In each case, AlexNet (trained on the ImageNet training set) assigned a confidence greater than 99 percent that the image was an instance of the category shown.

Imperfect Decision Making

Confusion matrix

	Actually Positive	Actually Negative
Called Positive	**True Positive**, power, **sensitivity**, recall, hit rate	**False Positive**, type I error, false alarm
Called Negative	**False Negative**, type II error	**True Negative**, **specificity**

Imperfect Models

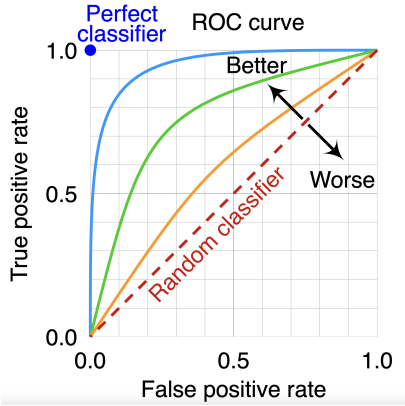

Figure 4: Wikipedia

By cmglee, MartinThoma, CC BY-SA 4.0

Less Imperfect Models [5]

AI-augmented clinical decision making

Australian Alliance for Artificial Intelligence in Healthcare

Friedman's Fundamental Theorem of Informatics (Friedman 2009)

"A person working in partnership with an information resource is 'better' than that same person unassisted"

Human + AI > Human > AI

Shift focus of AI tools from **end-to-end decision making** to ***supporting humans*** to make better clinical decisions

[5] Karin Verspoor

How to merge/partner with information resource is an open question (Cabitza et al. 2023)

- AI support was found useful but XAI was associated with a null or detrimental effect.
- AI-first protocols had higher accuracy than human-first ones and humans or AI alone.

How automated do we want medicine to be? (Topol 2019)

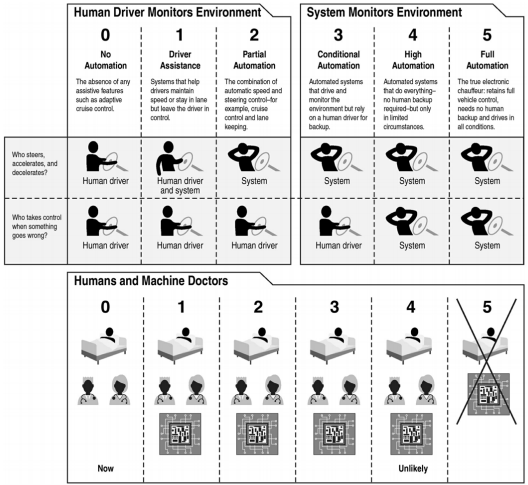

FIGURE 4.8: Self-driving cars and medicine. The Society of Automotive

Data Scarcity (Johnson et al. 2022)

The Ghost in the Machine has an American accent: value conflict in GPT-3.

Rebecca L Johnson
The University of Sydney, Australia.
Rebecca.johnson@sydney.edu.au

Giada Pistilli
Sorbonne Université, France.
Giada.pistilli@paris-sorbonne.fr

Natalia Menédez-González
European University Institute, Spain.
Natalia.menendez@eui.eu

Leslye Denisse Dias Duran
Ruhr Universität Bochum, Germany.
Leslye.diasduran@ruhr-uni-bochum.de

Enrico Panai
University of Sassari, Italy.
enricopanai@gmail.com

Julija Kalpokiene
Vytautas Magnus University, Lithuania.
Julija.kalpokiene@vdu.lt

Donald Jay Bertulfo
Delft University of Technology, Netherlands.
d.j.bertulfo@tudelft.nl

Section 5

Current Applications in Healthcare

Victoria, Australia

- "AI" is currently "banned" in healthcare
  - But Epic and Cerner have decision support tools that are "AI"
- And I'm looking at large language models for text simplification and summarization for patients

Intermountain Healthcare

- Pneumonia surveillance
  - Image processing
  - Text analysis

The most FDA approved “AI” tools are in imaging

- Radiology
- Ophthalmology

AI in Mammography (Jackson 2023)

Article • Data & Analytics

AI used to read breast cancer screenings a safe success

By **Amber Jackson**

August 03, 2023 • 4 mins

Lead author of the study Dr Kristina Lång, from Lund University in Sweden, said: "The greatest potential of AI right now is that it could allow radiologists to be less burdened by the excessive amount of reading. "While our AI-supported screening system requires at least one radiologist in charge of detection, it could potentially do away with the need for double reading of the majority of mammograms, easing the pressure on workloads and enabling radiologists to focus on more advanced diagnostics while shortening waiting times for patients."

AI in Mammography (Satariano, Metz, and Akos 2023)

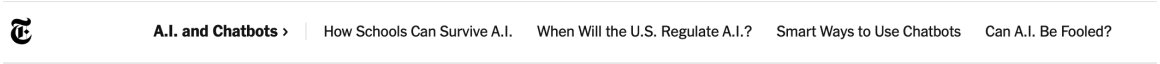

Using A.I. to Detect Breast Cancer That Doctors Miss

Hungary has become a major testing ground for A.I. software to spot cancer, as doctors debate whether the technology will replace them in medical jobs.

"An A.I.-plus-doctor should replace doctor alone, but an A.I. should not replace the doctor," Mr. Kecskemethy said.

The National Cancer Institute has estimated that about 20 percent of breast cancers are missed during screening mammograms.

Dr. Constance Lehman, a professor of radiology at Harvard Medical School and a breast imaging specialist at Massachusetts General Hospital, urged doctors to keep an open mind.

"We are not irrelevant," she said, "but there are tasks that are better done with computers."

Section 6

Ethical Issues

Some challenges of using AI (from Simon D'Alfonso)

- Lack of Transparency in AI Algorithms
- Challenges in Assigning Responsibility
  - Credit attributions in cases of AI success?
  - Blame attributions in cases of AI failure?
- Importance of Explainable AI (XAI)
  - Interpretability
  - Explainability

Cognitive Outsourcing (Danaher 2018)

As Selinger and Frischmann (2016) have recently noted, usage of AI assistance is effectively a new form of outsourcing. ***Humans have long outsourced the performance of cognitive tasks to others.*** *I don't do my tax returns; my accountant does. I don't book my travel arrangements; my assistant does. Such humanistic outsourcing has its own ethical issues. . . .Michael Sandel (2012) has argued, there are some tasks that seem to ethically demand my personal involvement. For instance, outsourcing the writing of a best man's speech seems like a mark of disrespect and apathy, not a praiseworthy efficiency-maximising way to fulfil one's duties.*

Cognitive Outsourcing (Danaher 2018)

One of those knock-on effects, according to Carr, ***is that increased reliance on AI assistance will atrophy and degenerate our mental faculties.*** *So, far from freeing up mental resources, increased reliance on AI assistance will deplete mental resources. We will no longer have the ability to think the important thoughts. This will in turn reduce the quality of our personal lives because the ability to engage in deep thinking is both intrinsically and instrumentally valuable: it results in a better immediate conscious experience and engagement with life, and it helps one to solve personal problems.*

Poll Everywhere

"The use of decision support systems in clinical education should be avoided as it will inevitably reduce the cognitive and diagnostic skills of the students."

Microsoft's responsible AI principles

1. Fairness - AI systems should treat all people fairly.
2. Reliability and safety - AI systems should perform reliably and safely.
3. Privacy and security - AI systems should be secure and respect privacy.
4. Inclusiveness - AI systems should empower everyone and engage people.
5. Transparency - AI systems should be understandable.
6. Accountability - People should be accountable for AI systems.

https://www.microsoft.com/en-us/ai/responsible-ai

“A Unified Framework of Five Principles for AI in Society” (Floridi and Cowls 2019)

1. Beneficence: Promoting Well-Being, Preserving Dignity, and Sustaining the Planet
2. Non-Maleficence: Privacy, Security and ‘Capability Caution’
3. Autonomy: The Power to Decide (to Decide)
4. Justice: Promoting Prosperity, Preserving Solidarity, Avoiding Unfairness
5. Explicability: Enabling the Other Principles through Intelligibility and Accountability

- epistemological sense of ‘intelligibility’ (as an answer to the question ‘how does it work?’)
- ethical sense of ‘accountability’ (as an answer to the question ‘who is responsible for the way it works?’)

Thank you

- Final attitudes towards AI in healthcare
- `PollEv.com/brianchapman270`
- `brianchapman270 to 22333`

References I

"Artificial Intelligence." 1980. *New England Journal of Medicine* 302 (26): 1481–82. https://doi.org/10.1056/NEJM198006263022617.

Braithwaite, J., P. Glasziou, and J. Westbrook. 2020. "The three numbers you need to know about healthcare: the 60-30-10 Challenge." *BMC Med* 18 (1): 102.

Bringsjord, Selmer, and Naveen Sundar Govindarajulu. 2022. "Artificial Intelligence." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman, Fall 2022. https://plato.stanford.edu/archives/fall2022/entries/artificial-intelligence/; Metaphysics Research Lab, Stanford University.

Cabitza, Federico, Andrea Campagner, Luca Ronzio, Matteo Cameli, Giulia Elena Mandoli, Maria Concetta Pastore, Luca Maria Sconfienza, Duarte Folgado, Marília Barandas, and Hugo Gamboa. 2023. "Rams, Hounds and White Boxes: Investigating Human–AI Collaboration Protocols in Medical Diagnosis." *Artificial Intelligence in Medicine* 138: 102506. https://doi.org/https://doi.org/10.1016/j.artmed.2023.102506.

References II

Clark, A. 2003. *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence. Oxford University Press. https://books.google.com.au/books?id=8JXaK3sREXQC.

D., LEE. 1977. "Experts Argue Whether Computers Could Reason, and If They Should: Experts Dispute Whether Computers Could Reason." *New York Times (1923-)*. https://www.proquest.com/historical-newspapers/experts-argue-whether-computers-could-reason-if/docview/123325923/se-2.

Danaher, John. 2018. "Toward an Ethics of Ai Assistants: An Initial Framework." *Philosophy and Technology* 31 (4): 629–53. https://doi.org/10.1007/s13347-018-0317-3.

Floridi, Luciano, and Josh Cowls. 2019. "A Unified Framework of Five Principles for AI in Society." *Harvard Data Science Review* 1 (1).

Fry, H. 2019. *Hello World: Being Human in the Age of Algorithms*. WW Norton. https://books.google.com.au/books?id=5l2NEAAAQBAJ.

References III

Jackson, Amber. 2023. “AI Used to Read Breast Cancer Screenings a Safe Success.” *AI Magazine*. https://aimagazine.com/articles/ai-used-to-read-breast-cancer-screenings-a-safe-success.

Johnson, Rebecca L, Giada Pistilli, Natalia Menédez-González, Leslye Denisse Dias Duran, Enrico Panai, Julija Kalpokiene, and Donald Jay Bertulfo. 2022. “The Ghost in the Machine Has an American Accent: Value Conflict in GPT-3.” https://arxiv.org/abs/2203.07785.

Medicine, I., L. A. M. Olsen, E. G. Nabel, J. M. McGinnis, and M. B. McClellan. 2008. *Evidence-Based Medicine and the Changing Nature of Health Care: 2007 IOM Annual Meeting Summary*. Learning Healthcare System Series. National Academies Press. https://books.google.com.au/books?id=xA5kAgAAQBAJ.

References IV

Mirsky, Yisroel, Tom Mahler, Ilan Shelef, and Yuval Elovici. 2019. "CT-GAN: Malicious Tampering of 3D Medical Imagery Using Deep Learning." In *Proceedings of the 28th USENIX Conference on Security Symposium*, 461–78. SEC'19. USA: USENIX Association.

Mitchell, M. 2019. *Artificial Intelligence: A Guide for Thinking Humans*. Farrar, Straus; Giroux. https://books.google.com.au/books?id=65iEDwAAQBAJ.

Satariano, Adam, Cade Metz, and Stiller F. Akos. 2023. "Using A.I. To Detect Breast Cancer That Doctors Miss." New York Times (Online). 2023. https://www.proquest.com/blogs-podcasts-websites/using-i-detect-breast-cancer-that-doctors-miss/docview/2782623735/se-2.

Shortliffe, E. H., S. G. Axline, B. G. Buchanan, T. C. Merigan, and S. N. Cohen. 1973. "An artificial intelligence program to advise physicians regarding antimicrobial therapy." *Comput Biomed Res* 6 (6): 544–60.

References V

Spackman, K. A. 1985. "A program for machine learning of counting criteria: empirical induction of logic-based classification rules." *Comput Methods Programs Biomed* 21 (3): 221–26.

Topol, E. 2019. *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. Basic Books. https://books.google.com.au/books?id=_EFlDwAAQBAJ.

WARNER, H. R., A. F. TORONTO, L. G. VEASEY, and R. STEPHENSON. 1961. "A mathematical approach to medical diagnosis. Application to congenital heart disease." *JAMA* 177 (July): 177–83.

Whitehead, A. N. 1958. *An Introduction to Mathematics*. A Galaxy Book, GB18. Oxford University Press. https://books.google.com.au/books?id=JmKM_0EZm5AC.