Pushing Against Ignorance: Medical Informatics, Artificial Intelligence, and the Quest to Improve Healthcare

Brian E. Chapman, PhD

2024-05-30

Why my title?

- Ignorance (and uncertainty) is a seemingly intrinsic (nearly defining) aspect of being a patient
- NLM G13 proposal: Pushing Against Ignorance: An Autoethnographic Study of How Informatics Has Shaped the Culture of Healthcare
- Moving to Australia has intensified my interest in “pushing against ignorance”

Ignorance

"Ignorance exists because man is a being of limited intelligence and power. . . .Accordingly, there is nothing incidental or fortuitous about our ignorance—it is something deep-rooted in the nature of things. (Rescher 2009, xi)

Yet “ignorance”. . .

- Doesn’t seem to be a common medical term (at least in my circles)
- Uncertainty, error—related, more common words

Ignorance (Rescher 2009)

- Uncertainty is the experience/awareness of ignorance
- Error may be the consequence of ignorance
- Invincible ignorance: that which is intrinsically unknowable or beyond the scope of current human knowledge
  - Basic science pushes against invincible ignorance
- “Vincible ignorance is that which an individual can overcome with a reasonable amount of effort.” (Rescher 2009, 12)
  - This is largely the domain of medical informatics

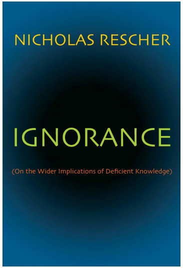

Section 1

Imaging and NLP

Alfred North Whitehead

"Plato and Pythagoras stand nearer to modern physical science than does Aristotle. . . .

"The popularity of Aristotelian Logic retarded the advance of physical science throughout the Middle Ages. If only the schoolmen had measured instead of classifying, how much they might have learnt!

"Classification is necessary. But unless you can progress from classification to mathematics, your reasoning will not take you very far."

~Alfred North Whitehead, *Science and the Modern World*

Quantitative imaging

- Anonymous chair of radiology circa 1995: “A radiologist with a ruler is a radiologist in trouble.”

Quantitative imaging

Carl Jaffe, M.D.: “No one in clinical [drug] trials takes radiology seriously.” (CaBIG, Dec. 2005)

- **“The inability to quantitatively monitor therapy”**

Non-quantitative radiological knowledge might be labeled “impoverished knowledge” (vincible ignorance)

Medical knowledge is “impoverished”

“Often it is not too much of a leap to infer [from a RCT] that the intervention causes the positive outcome.” **But the resulting knowledge is rather impoverished: it is knowledge of what works, without knowledge of how it works (or why it does not work), or how to make it work better. It is knowledge of effects without knowledge of underlying mechanisms.**” (Solomon 2015, 117)

Figure 1: Miriam Solomon, Wikipedia

Compare: black box machine learning

Quantitative Imaging: Vascular

Knowledge we can reason with

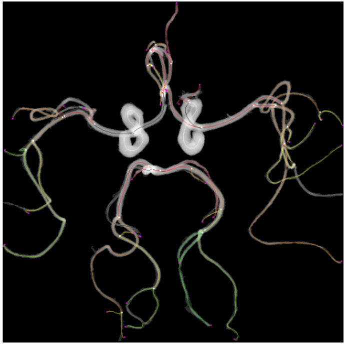

Great opportunities with deep learning

pyConTextNLP and scriptable NLP

- Transition from imaging to text
- Using dictated reports as labels for images
  - Uncertainty
  - Negation
  - Location

pyConTextNLP and scriptable NLP

- Amil Gentili, M.D.
- How to help Amil do critical finding reporting auditing?
- Help him pick up basic software/programming skill

Section 2

Educational efforts

The educational challenge of biomedical informatics

through an electrical engineering metaphor

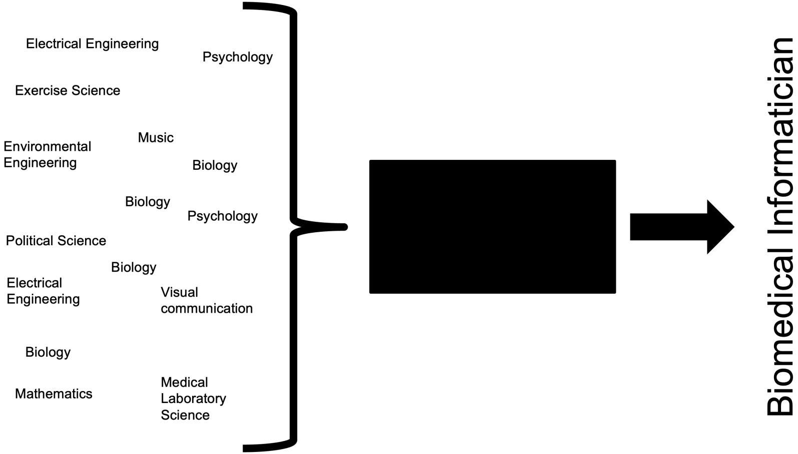

The educational challenge of biomedical data science

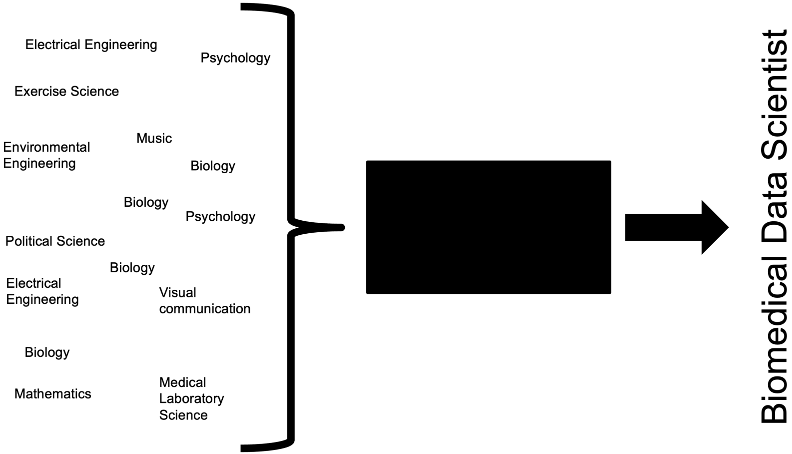

Another metaphor [1]

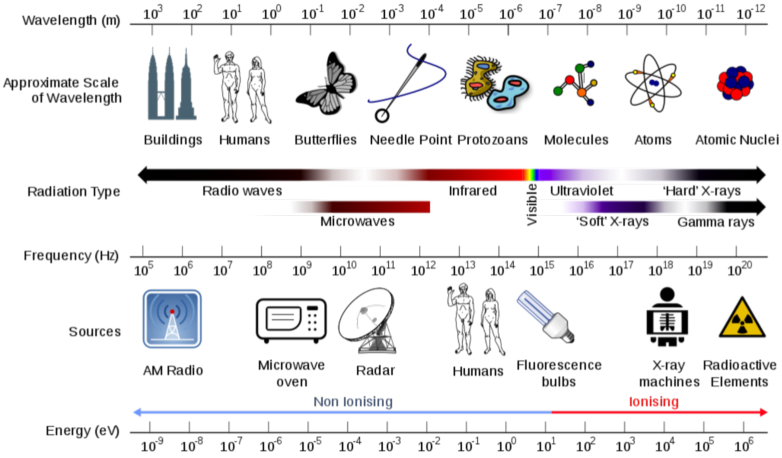

[1] Wikipedia

Is everything the same along the spectrum?

- It is all electromagnetics!
- But...
  - Low Frequencies I might ignore wave properties (Kirchhoff's Current and Voltage Laws)
  - At High Frequencies I might also ignore wave properties (photons)
- Different skill sets and tools needed to be productive
- Yet it is all electromagnetics!

Another metaphor [2]

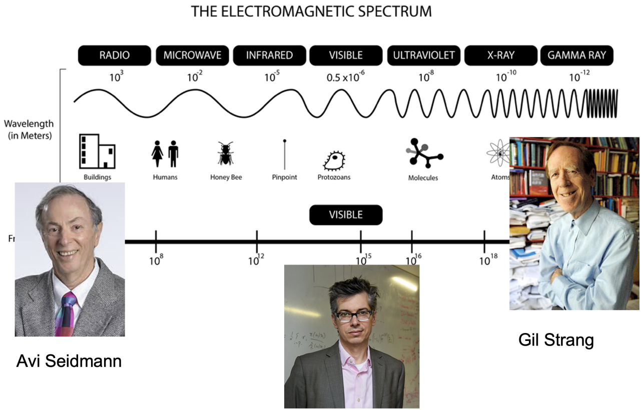

Avi Seidmann

Chris Wiggins

Gil Strang

[2] Wikipedia

Punk Data Science

Utah R25: “Data, exploration, Computation, and Analytics Real-world Training (DeCART)”

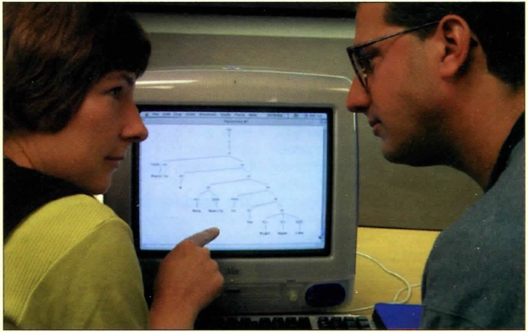

Wendy Chapman (left) and Marcelo Fiszman, M.D., both Ph.D. candidates in medical informatics, discuss the parsed results of their work in “natural language processing.” Future use of natural language processing techniques will have high utility in the development of enterprise clinical information systems.

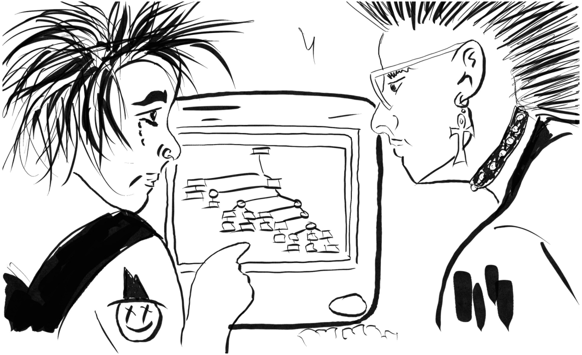

Punk Informatics: Homer Warner, MD, PhD

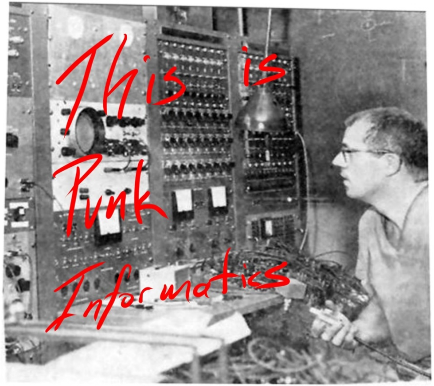

Punk Data science

"[P]unk is not a narrow musical style or a particular fashion or hairstyle. Rather, it is a commitment to a **DIY sensibility and, with that, a dedication to self-empowerment.**" (Dunn 2016, 15)

Punk Data science

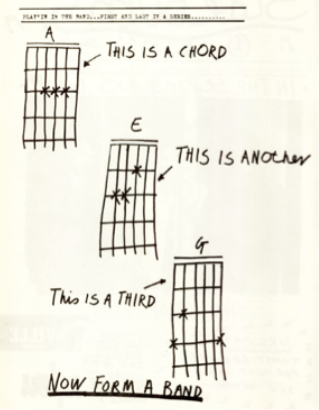

Section 3

Artificial Intelligence

First Melbourne tasks

- Create data literacy
  - Punk data science!
- Biomedical informatics and artificial intelligence
- All with 5 hours per year of student contact time!

How to narrow the task?

- Focus on physician as a decision maker
- Not difficult to motivate AI from this perspective

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

Decision making in medicine

Overwhelming amount of information (Medicine et al. 2008)

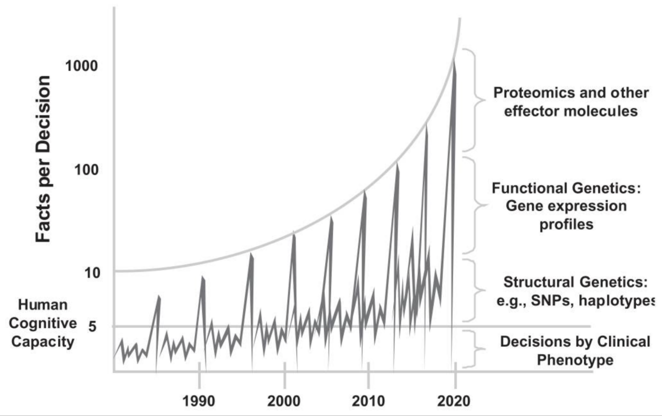

What are brains good for? Practicing medicine?

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)


[a]

[a] The New Yorker

Human brain exceptionalism

[W]hat 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
  - paper and pencil to do arithmetic
  - books for storing and recalling knowledge
  - Cell phones, computers
- Extended mind theory

"Man-Computer Symbiosis" (Licklider 1960)

4 IRE TRANSACTIONS ON HUMAN FACTORS IN ELECTRONICS March

Man-Computer Symbiosis*

J. C. R. LICKLIDER†

Summary—Man-computer symbiosis is an expected development in cooperative interaction between men and electronic computers. It will involve very close coupling between the human and the electronic members of the partnership. The main aims are 1) to let computers facilitate formulative thinking as they now facilitate the solution of formulated problems, and 2) to enable men and computers to cooperate in making decisions and controlling complex situations without inflexible dependence on predetermined programs. In the anticipated symbiotic partnership, men will set the goals, formulate the hypotheses, determine the criteria, and perform the evaluations. Computing machines will do the routinizable work that must be done to prepare the way for insights and decisions in technical and scientific thinking. Preliminary analyses indicate that the symbiotic partnership will perform intellectual operations much

will be coupled together very tightly, and that the resulting partnership will think as no human brain has ever thought and process data in a way not approached by the information-handling machines we know today.

B. Between "Mechanically Extended Man" and "Artificial Intelligence"

As a concept, man-computer symbiosis is different in an important way from what North[2] has called "mechanically extended man." In the man-machine systems of the past, the human operator supplied the initiative, the direction, the integration, and the criterion. The

"The hope is that, in not too many years, human brains and computing machines will be coupled together very tightly and that the resulting partnership will think as no human brains as ever thought. . . ."

Example: Decision Support (circa 1970)

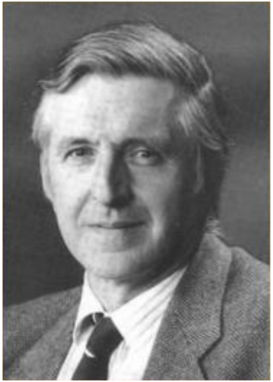

Figure 2: FT de Dombal

- AAPHelp: Bayesian diagnosis of abdominal pain

Example: Decision Support (circa 1970)

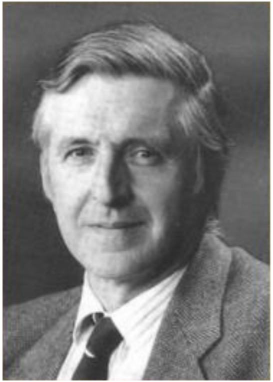

Figure 2: FT de Dombal

- AAPHelp: Bayesian diagnosis of abdominal pain
- “de Dombal’s team collected data on thousands of patients who presented with acute abdominal pain. The researchers used data on clinical symptoms (e.g., pain severity, location, and quality) and physical signs (e.g., pulse and abdominal guarding) to derive probabilities for the computer system.”

Example: Decision Support (circa 1970)

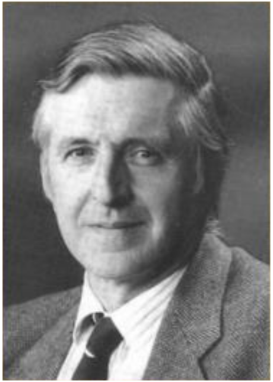

Figure 2: FT de Dombal

- AAPHelp: Bayesian diagnosis of abdominal pain
- "de Dombal's team collected data on thousands of patients who presented with acute abdominal pain. The researchers used data on clinical symptoms (e.g., pain severity, location, and quality) and physical signs (e.g., pulse and abdominal guarding) to derive probabilities for the computer system."
- Evaluated from 1971-1972. Outperformed senior clinicians with a correct diagnosis 91.8% of the time

Example: Decision Support (circa 1970)

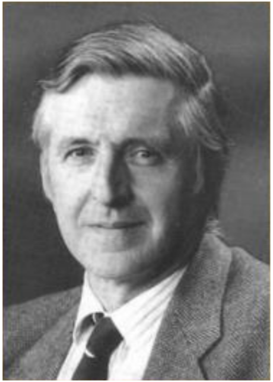

Figure 2: FT de Dombal

- AAPHelp: Bayesian diagnosis of abdominal pain
- “de Dombal’s team collected data on thousands of patients who presented with acute abdominal pain. The researchers used data on clinical symptoms (e.g., pain severity, location, and quality) and physical signs (e.g., pulse and abdominal guarding) to derive probabilities for the computer system.”
- Evaluated from 1971-1972. Outperformed senior clinicians with a correct diagnosis 91.8% of the time
- Amazing! This should be exported!

Example: Decision Support (circa 1970)

- "But when his group teamed up with researchers at Bispebjerg Hospital in Copenhagen in 1976 to test the system in a fresh clinical environment, its overall accuracy plummeted to 65%."
- **Note: Similar results when Cerner recently tried to port their sepsis detector from USA to Australia**

What is going on?

- “databases don’t travel” (de Dombal)

What is going on?

- “databases don’t travel” (de Dombal)
- **“clinical data are formed, not found.”** (Lea and Jones 2024)

What is going on?

- “databases don’t travel” (de Dombal)
- **“clinical data are formed, not found.”** (Lea and Jones 2024)
- Healthcare is culturally dependent and context sensitive.

What is going on?

- "databases don't travel" (de Dombal)
- **"clinical data are formed, not found."** (Lea and Jones 2024)
- Healthcare is culturally dependent and context sensitive.
- Question: how much generalizable knowledge can be created in healthcare (as opposed to physiology, molecular biology, etc.)?

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

What are the implications for the development and use of AI in the USA (population ~333M) vs. Australia (population ~26M)?

Implications of data-driven AI (Braithwaite, Glasziou, and Westbrook 2020)

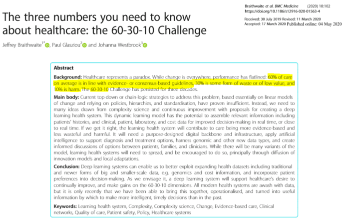

Braithwaite *et al. BMC Medicine* (2020) 18:102
https://doi.org/10.1186/s12916-020-01563-4
Received: 30 July 2019 Revised: 11 March 2020
Accepted: 17 March 2020 Published online: 04 May 2020

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

These suboptimal processes will create suboptimal data. Should we be building AI tools based on these data?

Improved mechanistic reasoning: Glucose management (Morris et al. 2021)

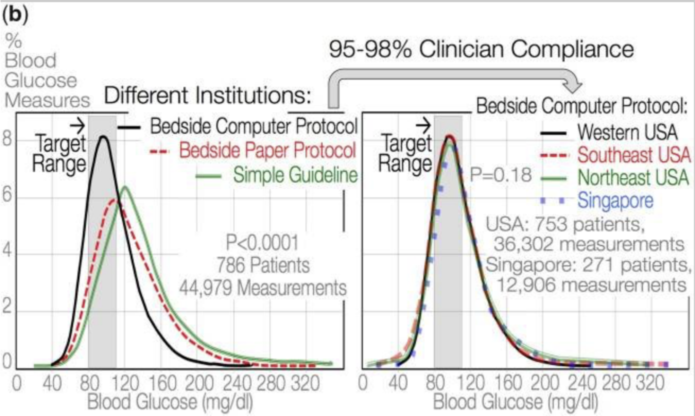

Example of AI for low resource countries

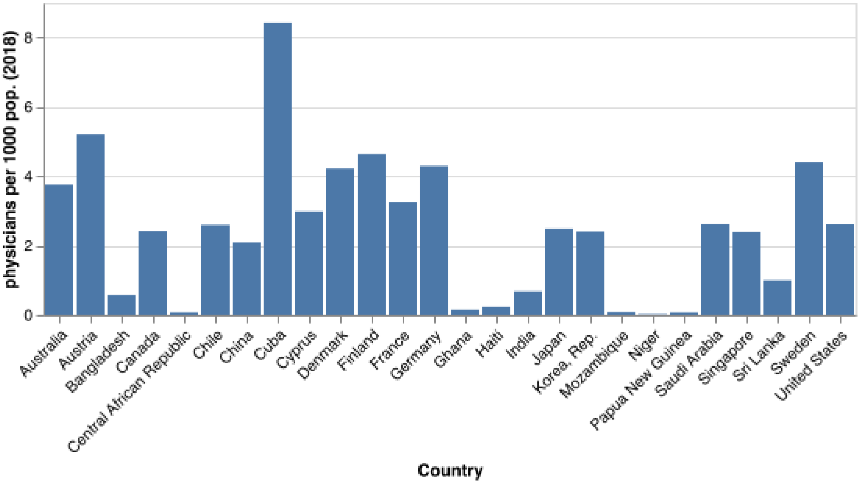

How doctors are using AI to diagnose a hidden heart condition in kids

The Washington Post
Democracy Dies in Darkness

How doctors are using AI to diagnose a hidden heart condition in kids

By Mark Johnson
January 16, 2024 at 5:00 a.m. EST

A handheld echo device and ultrasound probe, right, sit next to the portable AI device developed by Children's National and Us2.ai during its initial testing phase in northern Uganda. (Children's National Hospital)

"Now, in an advance that shows the potential of artificial intelligence to aid medicine, researchers at Children's National have developed a new AI-powered tool for diagnosing rheumatic heart disease long before a patient needs surgery. Collaborating with staff at the Uganda Heart Institute, the team designed a system that will allow trained nurses to screen and diagnose children early on, when they can still be treated with penicillin for less than $1 a year. Early treatment could save thousands from having to undergo surgery." (See (Brown et al. 2024) for scientific paper.)

Let's summarize this

- AI is allowing less skilled clinicians to accurately obtain clinical data

Let's summarize this

- AI is allowing less skilled clinicians to accurately obtain clinical data
- How far can we generalize this?

Let's summarize this

- AI is allowing less skilled clinicians to accurately obtain clinical data
- How far can we generalize this?
- **How can AI help patients/consumers/citizens better pursue and achieve their health goals?**

Section 4

AI and the patient experience

Epistemic injustice: Miranda Fricker

- Testimonial injustice
- **Hermeneutical injustice**
- **How can AI be used to address these injustices?**

Review some of the context of the patient experience

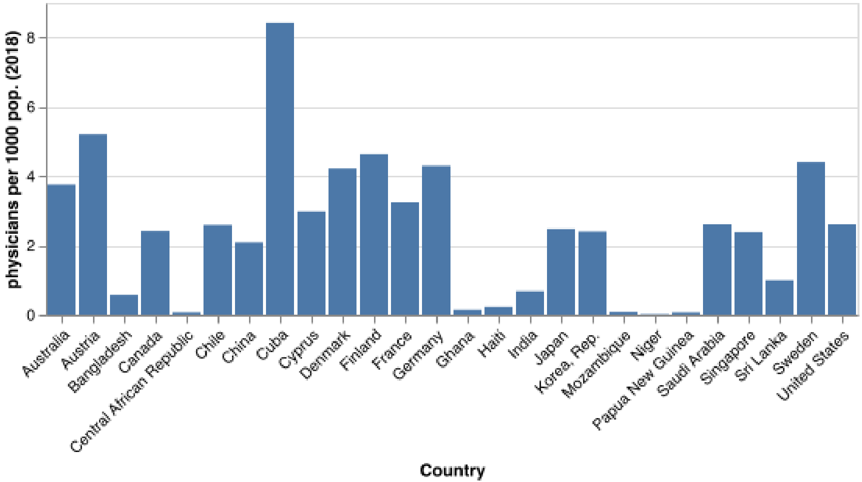

Context: Healthcare occurs in time-constrained settings (Irving et al. 2017)

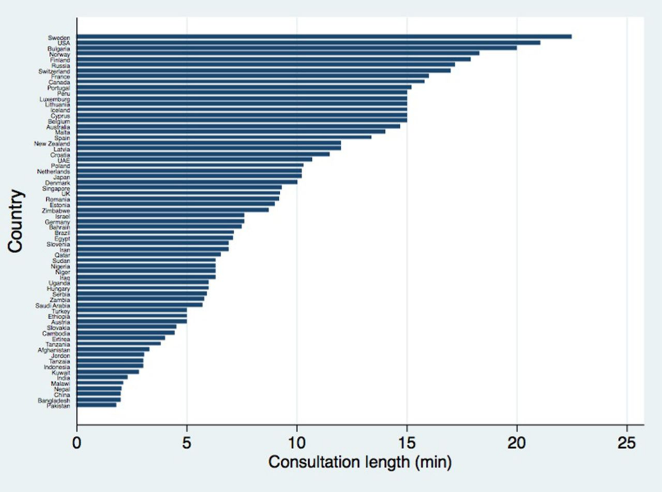

Context: Healthcare occurs in time-constrained settings (Irving et al. 2017)

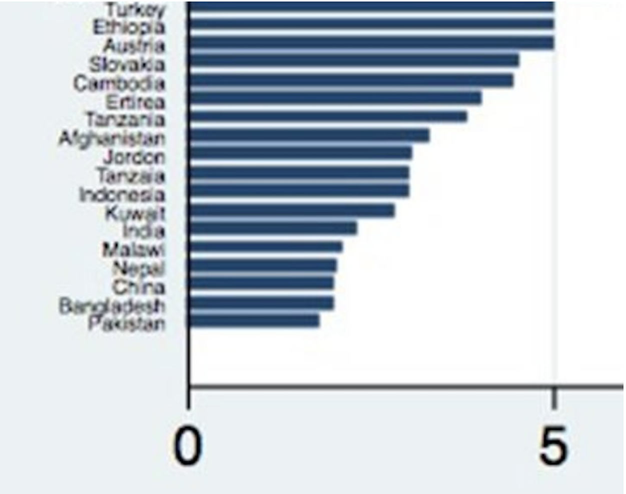

Context: healthcare occurs in noisy, error-prone environments

"We are error-prone and error tolerant—errors are unavoidable in the fabric of our lives. . . . Cognitively speaking, we metabolize mistakes!"

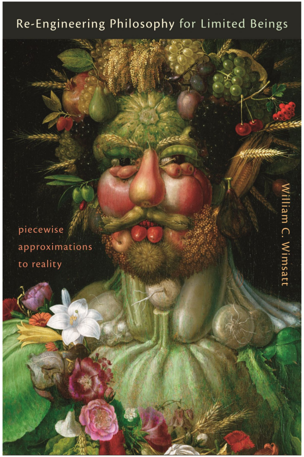

Our engineered solutions also need to be able to "metabolize mistakes"

Hermeneutical Injustice and AI

- Patients have a right to understand their health and their healthcare
- Hermeneutical injustice results when patients experience **vincible ignorance**
  - Not due to invincible ignorance
- Substantial engineered barriers to understanding
- LLMs might be beneficial to addressing some of these barriers

Ignorance is not a stable state

“Jumping to conclusions over a chasm of ignorance is a natural human tendency from which few of us are exempt.” (Rescher 2009, 2)

“Nature abhors a vacuum. So does the human mind. We try not to let the gaps in our knowledge be mere empty blanks, so we fill them in with speculation and suppositions.” (Rescher 2009, 14)

Can AI help patients fill in the gaps better? Can AI help create punk patients?

LLMs and vincible ignorance: With Sterre, Edward, and Maolin

- Text simplification, summarization, and translation works remarkably well
  - Constrained problem!
- What about question answering?
  - Very open/unconstrained problem!



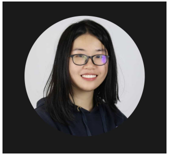

What about question answering?

- LLMs fairly good at answering, but…
  - Hallucination
  - And many more issues
- How to constrain answer?
- How do we help the patient “metabolize mistakes” from LLM?

Current exploration: RAG

- Constrain system to only answer questions with specified material
  - Tell the patient where that material came from
- Retrieve materials from reliable sources (PubMed, PubMed Central, NLM Bookshelf, etc.)
  - **Specificity of retrieval is the pain point**

Patient question answer systems

- Resource constrained solutions
  - The creativity of poverty
- Novel PubMed searching
  - Fine tuned embeddings
  - MeSH
  - Co-author graphs
- Subset text to pass only relevant text to LLM
  - ColBERT

Patient question answer systems

MetaGPT: The Multi-Agent Framework

Assign different roles to GPTs to form a collaborative entity for complex tasks.

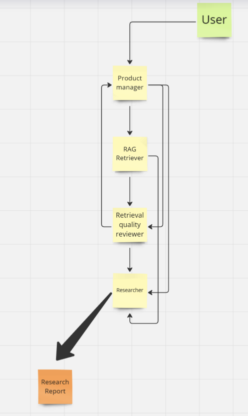

Next steps?

- Constrain system with explicit knowledge representation

Next steps?

- Constrain system with explicit knowledge representation
  - Not really sure how but convinced this is important

Next steps?

- Constrain system with explicit knowledge representation
  - Not really sure how but convinced this is important
- View LLM as a component in a larger circuit

Next steps?

- Constrain system with explicit knowledge representation
  - Not really sure how but convinced this is important
- View LLM as a component in a larger circuit
- Efficient and realistic evaluations (on a budget)

Next steps?

- Constrain system with explicit knowledge representation
  - Not really sure how but convinced this is important
- View LLM as a component in a larger circuit
- Efficient and realistic evaluations (on a budget)
  - Harry Barrett and the dogma that quality must always be related to specific tasks!

Summary

- Ignorance is central to the human experience, including healthcare
- Some ignorance is vincible
- Medical informatics pushes against the vincible ignorance in healthcare
  - Quantification
  - Mechanistic thinking
  - AI
- Clinicians and patients can partner with and benefit from this effort
- Punk (DIY, dedication to self-empowerment) gets easier with each generation of AI

Thank you

References I

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.

Brown, K., P. Roshanitabrizi, J. Rwebembera, E. Okello, A. Beaton, M. G. Linguraru, and C. A. Sable. 2024. “Using Artificial Intelligence for Rheumatic Heart Disease Detection by Echocardiography: Focus on Mitral Regurgitation.” *J Am Heart Assoc* 13 (2): e031257.

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.

Dunn, K. 2016. *Global Punk: Resistance and Rebellion in Everyday Life*. Bloomsbury Academic. https://books.google.com/books?id=jfdaEAAAQBAJ.

References II

Irving, G., A. L. Neves, H. Dambha-Miller, A. Oishi, H. Tagashira, A. Verho, and J. Holden. 2017. "International variations in primary care physician consultation time: a systematic review of 67 countries." *BMJ Open* 7 (10): e017902.

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.

Lea, A. S., and D. S. Jones. 2024. "Mind the Gap - Machine Learning, Dataset Shift, and History in the Age of Clinical Algorithms." *N Engl J Med* 390 (4): 293–95.

Licklider, J. C. R. 1960. "Man-Computer Symbiosis." *IRE Transactions on Human Factors in Electronics* HFE-1 (1): 4–11. https://doi.org/10.1109/THFE2.1960.4503259.

References III

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.

Morris, A. H., B. Stagg, M. Lanspa, J. Orme, T. P. Clemmer, L. K. Weaver, F. Thomas, et al. 2021. "Enabling a learning healthcare system with automated computer protocols that produce replicable and personalized clinician actions." *J Am Med Inform Assoc* 28 (6): 1330–44.

Rescher, N. 2009. *Ignorance: (On the Wider Implications of Deficient Knowledge)*. University of Pittsburgh Press. https://books.google.com/books?id=g0bw4or_iaQC.

Solomon, M. 2015. *Making Medical Knowledge*. Oxford University Press. https://books.google.com/books?id=___dsBwAAQBAJ.