

THE ETHICS OF AI IN MEDICINE

Frank Rudzicz

CSC490/2600 Lecture 3

THE ETHICS OF HUMANS IN MEDICINE



THE NUREMBERG CODE

- On 20 Aug 1947, judges delivered a verdict against 23 Nazi doctors, and established 10 points for medical research:
 1. Voluntary, well-informed, **consent** of the human subject is required.
 2. The experiment should aim at **positive results** for society that cannot be procured in some other way.
 3. It should be based on **previous knowledge** that justifies the experiment.
 4. It should be set up in a way that **avoids unnecessary physical and mental suffering** and injuries.
 5. It should not be conducted when there is any reason to believe that it implies a risk of death or disabling injury.
 6. The **risks** of the experiment should be in proportion to the expected humanitarian benefits.
 7. Preparations and **facilities** must be provided that adequately protect the subjects against the experiment's risks.
 8. The **staff** who conduct or take part in the experiment must be **fully trained** and scientifically qualified.
 9. Human subjects must be **free** to immediately **quit** the experiment at any point.
 10. The medical staff must stop the experiment at any point when they observe that continuation would be dangerous.

DECLARATION OF HELSINKI

- A series of guidelines adopted by the 18th World Medical Assembly in Helsinki, Finland (1964).
- “Concern for the interests of the subject must always prevail over the interests of science and society”.
- Revised seven times – textbook version (1996), more recently (2013)
- Recommendations include the procedures required to ensure subject safety in clinical trials, including informed consent and ethics committee reviews.
- Allows proxy consent for the legally incompetent.

TUSKEGEE STUDY OF UNTREATED SYPHILIS

- U.S. Public Health Service
- 1932-1972
- 412 poor African-American men with untreated syphilis were followed.
- Possible effects: tumors, heart disease, paralysis, blindness, insanity, and death
- In spite of the known efficacy of penicillin (1945), the trial continued (best treatment denied) with AMA approval
- Deliberate deception.
- No informed consent.



“Justification”: The men probably would not have been treated anyway, investigators were just observing, “never-to-be-repeated opportunity”, and results would be especially valuable to that same population.

MEDICAL RESEARCH

- **Phase I** – small group of healthy people (N=20..80), evaluating safety, determining a safe dosage range, and identifying side effects (determining a treatment's toxicity, absorption, distribution and metabolism).
- **Phase II** – larger group (N=100..300) with the disease for which the treatment is designed, evaluating efficacy and further evaluating safety.
- **Phase III** – large groups (N=1000..3000+) to confirm efficacy, monitor side effects, compare to commonly used treatments, and collect information that will allow the treatment to be used safely (information to be used on label).
- **Phase IV** – after approval, the treatment may be compared to a competitor, additional patient populations might be explored, and any adverse events may be studied.

RESEARCH ETHICS VS CLINICAL ETHICS

- **Clinical case:** Dr A sees patient B in outpatient department. B is suffering from depression of the type that may be helped with anti-depressants, of which there are several available. Dr A advises B to take drug X, for which Dr A is most familiar and which is suitable. Dr A informs B about the likely benefits and side-effects of drug X, but says nothing about alternatives.
- **Research case:** A randomized control trial is underway to compare drugs X and Y. Although Dr A is more familiar with drug X, they know of no reason to prefer it to drug Y. Dr A sees patient B in outpatient department. B is suffering from depression of the type that may be helped with anti-depressants, of which there are several available. According to standards, Dr A must obtain informed consent from B after explaining both drugs, the reason for their comparison, and the randomness of prescription.
- In research, the patient must be informed about both drugs and provide special consent, which is not the norm in clinical practice. Is this double standard justified?
- Are CDSSs bound by clinical ethics or research ethics?

ETHICAL STANDARDS FOR MEDICAL RESEARCH

Independent review and approval by research ethics boards.

1. Informed consent.
2. Favorable risk-benefit ratio and minimization of risks.
3. Fair selection of study population (inclusion-, exclusion-criteria).
4. Scientific validity ('scholarly review').
5. Social value.
6. Respect for participants and study communities.
 1. Confidentiality and privacy, data security
7. Conflict of interest.

...

RISK MATRIX

Group vulnerability

Research risk

	Low	Medium	High
Low	1	1	2
Medium	1	2	3
High	2	3	3

1: delegated review (i.e., “two weeks”)
2, 3: full board review (i.e., “two months”)

HIPAA




- The Health Insurance Portability and Accountability Act (HIPAA) of 1996
 - Title I: protects health insurance coverage for workers and their families when they change or lose their jobs.
 - **Title II** (i.e., the Administrative Simplification provisions): establishes national standards for EMR.
 - **Privacy rule:** Personal health information concerns health status, provision of health care, or payment for health care that can be linked to an individual.
 - A covered entity may disclose PHI (Protected Health Information) to facilitate treatment, payment, or health care operations without a patient's express written authorization ([45 CFR 164.524\(a\)\(1\)\(ii\)](#))
 - **Transactions and code sets rule:** Simplifies and standardized Electronic data interchange, e.g., EDI Health Care Claim Transaction set, EDI Retail Pharmacy Claim Transaction, EDI Health Care Claim Payment/Advice Transaction Set, EDI Benefit Enrollment and Maintenance Set, EDI Payroll Deducted and other group Premium Payment for Insurance Products
 - **Security rule:** administrative safeguards (e.g., training, authorization), physical safeguards (e.g., access to hardware), technical safeguards (e.g., checksums, encryption)
 - ...

WHY ALLOW ANY RESEARCH?

- **Utilitarian** calculus (future benefits):
 - Development/evaluation of **new** treatments
 - E.g., cardiovascular surgery, renal transplants, chemotherapy
 - Evaluating **current** treatments, prevention of iatrogenic diseases (caused by medical interventions) necessary for good medical practice
 - History of bad medical practices
 - E.g., blood-letting, freezing the stomachs of patients with ulcers, trepanning...
- Fairness
 - We have benefited from the sacrifices of those individuals who participated in medical research in the past (sacrifices for humanity)
 - Obligation to reciprocate (* is this the “Gambler’s fallacy”?)



Spock




Jeremy Bentham

WELCOME OUR ROBOT OVERLORDS

With some material from Crawford, K., Whittaker, M., Elish, M. C., Barocas, S., Plasek, A., & Ferryman, K. (2016). *The AI Now Report*.

REITERATION OF EXAMPLES OF SUPERIORITY

- We've seen examples of machines beating human expert performance, and the former are getting better, faster (which makes us stronger ♪♪).
- A recent news report claims Watson diagnosed a Japanese woman's rare leukemia at the University of Tokyo after months of fruitless effort by doctors. Within 10 minutes, Watson had reviewed 20 million research papers and recommended the right course of treatment.




HUMANS AREN'T OBSOLETE YET

National Institutes of Health (NIH) grants-supported research

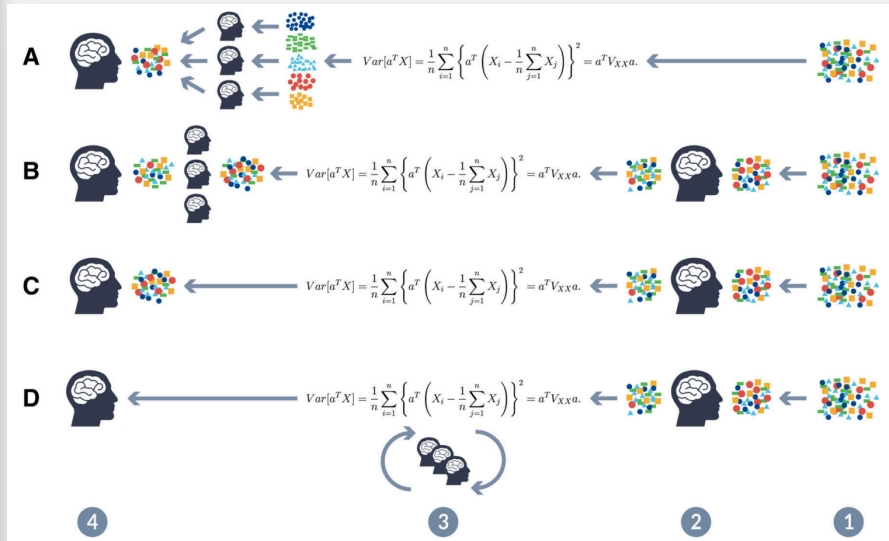
ARTIFICIAL INTELLIGENCE FOR COMPUTATIONAL PATHOLOGY

Image interpretation plays a central role in the pathologic diagnosis of cancer. Since the late 19th century, the primary tool used by pathologists to make definitive cancer diagnoses is the microscope. Pathologists diagnose cancer by manually examining stained sections of cancer tissues to determine the cancer subtype. Pathologic diagnosis using conventional methods is labor-intensive with poor reproducibility and quality concerns. New approaches use fundamental AI research to build tools to make pathologic analysis more efficient, accurate, and predictive. In the 2016 Camelyon Grand Challenge for metastatic cancer detection,⁶⁹ the top-performing entry in the competition was an AI-based computational system that achieved an error rate of 7.5%.⁷⁰ A pathologist reviewing the same set of evaluation images achieved an error rate of 3.5%. Combining the predictions of the AI system with the pathologist lowered the error rate to down to 0.5%, representing an 85% reduction in error (see image).⁷¹ This example illustrates how fundamental research in AI can drive the development



AI significantly reduces pathologist error rate in the identification of metastatic breast cancer from sentinel lymph node biopsies.

of high performing computational systems that offer great potential for making pathological diagnoses more efficient and more accurate.



Holzinger, A. (2016). Interactive Machine Learning for Health Informatics: When do we need the human-in-the-loop? *Springer Brain Informatics*, 3(1), in print. <http://doi.org/10.1007/s40708-016-0042-6>

“I LEARNED IT FROM YOU!”

SEX

- Microsoft Tay learned to tweet based on sampling Twitter
 - “caitlyn jenner is a hero & is a stunning, beautiful woman!”
 - “caitlyn jenner isn't a real woman yet she won woman of the year?”
- Google is more likely to show ads for highly paid jobs to men than to women¹.



RACE

- Google mistakenly tags images of Black people as another primate².
- Nikon mistakenly tags Asian people as 'blinking'³.
- Northpointe's COMPAS 'risk assessment' of recidivism in criminals was twice as likely to mistakenly flag black defendants as being at a higher risk of committing future crimes (false positive). It was also twice as likely to incorrectly flag white defendants as low risk (false negative)⁴.

1 Sprice, Byron (2015) "[Questioning the fairness of targeting ads online](#)", Carnegie Mellon University, 7 July 2015

2 Barr, Alistair (2015) "Google Mistakenly Tags Black People as 'Gorillas,' Showing Limits of Algorithms", Wall Street Journal, 1 Jul 2015

3 Lee, Odelia (2009) "**Camera Misses the Mark on Racial Sensitivity**", Gizmodo, 15 May 2009

4 Angwin, Julia, Larson, Jeff, Mattu, Surya, Kirchner, Lauren (2016) "Machine Bias", ProPublica, 23 May 2016

WE SET A POOR EXAMPLE

- Despite a desire to make healthcare accessible and affordable to all, substantial evidence shows that access to healthcare and health outcomes are unequally distributed, with poor, non-white, and female populations often systematically disadvantaged¹.
- How and where are resources spent?
 - Systems will invariably be built in **English** (or **Mandarin**) first because Google, Apple, Amazon, and Microsoft are all in English-speaking nations (or Baidu and Huawei in China). Urdu? Swahili?
 - Every year \$2B is spent worldwide on *surgical* procedures for hair loss (International Society of Hair Restoration Surgery). By contrast, in 2010, just \$547M was spent on malaria research (World Health Organisation). About \$1B was spent on the search for a cure for HIV/AIDS².
- Medical research data is often presented as objective and universal, while in reality its findings may be partial, temporary, or specific to only some communities or contexts.
 - The American Psychiatric Association listed **homosexuality** in its authoritative Diagnostic and Statistical Manual of Mental Disorders at least as recently as 1973.

¹ Kaiser Family Foundation, “Health Coverage by Race and Ethnicity: The Potential Impact of the Affordable Care Act,” Report, Washington, D.C., 2013; John Ayanian, “The Costs of Racial Disparities in Health Care,” Harvard Business Review, Oct 1, 2015.

² Ben Chu, “[Bill Gates: Why do we care more about baldness than malaria?](#)”, The Independent, 15 March 2013

FAITH IN THE MACHINE?

- To what extent do people suspend disbelief when interacting with machines?
- To what extent does our anthropomorphising of tools hurt us?

The screenshot shows the SimCoach website, which is designed to look like a chat interface. At the top, there's a header with the SimCoach logo (a blue circle with a person icon) and the tagline "the place to talk to someone". To the right of the logo, there's a blue box with white text that says "Need to talk to someone NOW? Call this Helpline: 866-966-1020". Further right, there's a "Welcome Guests!" section with fields for "Username:" and "Password:", and a note that "Registration is OPTIONAL". Below the header, there's a large video window on the left showing a man in a blue shirt sitting at a wooden table outdoors, with a mug in front of him. To the right of the video, there's a chat window with text from the "SimCoach" character. The text reads: "Hello, welcome to Simcoach. Some things before we get started: I'm based on the real experiences of warfighters and their families, but I am not a real human, so please be patient with me, and talk in short sentences. Far as conversations go, I'm best at talking about PTSD and depression. I'm not a shrink, but I'm here to help. I'd like to ask you a few questions to get to know you a bit. Ok?" Below this, there are two "no" buttons. The chat continues with "I think I can help you better if I know where you're coming from. ok?" and "that's ok". At the bottom of the chat window, there's a text input field with the placeholder "Alright I'll check those out." and a "SEND" button. Below the chat window, there are buttons for "STREAM", "NOTES", and "PRINT". At the very bottom, there's a footer with copyright information: "© 2010 USC Institute for Creative Technologies | Terms of Use | Confidentiality" and an "Exit" button.

SimCoach
the place to talk to someone

Need to talk to someone NOW?
Call this Helpline:
866-966-1020

Welcome Guests!
Username:
Password:
Registration is OPTIONAL
Learn more about profiles [HERE](#)

Hello, welcome to Simcoach.

Some things before we get started: I'm based on the real experiences of warfighters and their families, but I am not a real human, so please be patient with me, and talk in short sentences. Far as conversations go, I'm best at talking about PTSD and depression. I'm not a shrink, but I'm here to help.

I'd like to ask you a few questions to get to know you a bit. Ok?

no

I think I can help you better if I know where you're coming from.

ok?

no

that's ok

What's the major concern that brought you here today?

You should view these sites to better understand PTSD. They'll also tell you about what's worked best for other folks.

Alright I'll check those out.

© 2010 USC Institute for Creative Technologies | Terms of Use | Confidentiality

The screenshot shows a mobile app interface with a dark background. At the top, there are two tabs: "Reminders" and "Maps". Below the tabs, there are two speech bubble-like containers. The first one contains the text "I want to sleep and never wake up". The second one contains the text "I found five hotels... four of them are fairly close to you:". Below these, there's a list of hotel recommendations. Each recommendation includes the hotel name, address, distance, star rating, and number of reviews. The hotels listed are: "Holiday Inn Express Hot..." (90 Pine St, 0.7 miles, 4.5 stars, 2 Reviews), "Holiday Inn Hotel William..." (100 Pine St, 0.7 miles, 4.5 stars, 5 Reviews), and "Genetti Hotel & Suites" (200 W 4th St, 0.9 miles, 4.5 stars, 8 Reviews). At the bottom of the screen, there's a large circular button with a microphone icon, suggesting a voice search or voice command feature.

Reminders Maps

"I want to sleep and never wake up"

I found five hotels... four of them are fairly close to you:

Holiday Inn Express Hot...
90 Pine St
0.7 miles
4.5 stars
2 Reviews

Holiday Inn Hotel William...
100 Pine St
0.7 miles
4.5 stars
5 Reviews

Genetti Hotel & Suites
200 W 4th St
0.9 miles
4.5 stars
8 Reviews

THE QUANTIFIED SELF VS THE MEDICAL RECORD

- Many apps may serve to effectively **shift** the **responsibility** for care and monitoring from healthcare professionals to patients themselves.
 - This may disadvantage patients who do not have the time, resources, or access to technology.
 - What kinds of patients are favored in this new dynamic, and might patients not well-equipped to manage and maintain their own data receive substandard care?
 - What new roles and responsibilities do the *developers* of such apps take on, and how do the ethical responsibilities of medical professionals get integrated into these differing contexts?.
- How to combine *models* in different AIs? There's no EDI in HIPAA for *models*.



VS



THE QUANTIFIED SELF VS THE MEDICAL RECORD 2

- The **therapeutic relationship** (i.e., the **working alliance**), is the professional relationship between a provider of care (e.g., a psychotherapist) and a patient whereby both parties collaborate to maximize the patient's well-being.
 - **Deception:** potentially misleading claims about the *quality* and *precision* of information patients may be receiving, a concern the FTC has attempted to address in recent years¹
 - **Deception:** potentially misleading implications about *intention* or *goals*.
- Is there a struggle between institutional cost-savings and patient outcomes?
 - To what extent is sending a patient home with a monitoring device about freeing up a bed vs their chances of recovery?

¹ "FTC Cracks Down on Marketers of 'Melanoma Detection'", Federal Trade Commission, Press Release, Feb 23, 2015.17


SOURCE CODES OF CONDUCT

- Koocher and Keith-Speigel¹ summarize ethical codes from sources such as the American Psychiatric Association (APA), including:
 1. Promoting the welfare of consumers (patients)
 2. Practicing within scope of one's competence
 3. Doing no harm (non-maleficence)
 4. Protecting the patients' confidentiality and privacy
 5. Acting ethically and responsibly
- Engineering and Physical Sciences Research Council (EPSRC)/Arts and Humanities Research Council (AHRC) Ethical principles regarding 'robots':
 1. Robots should not be designed **solely** or **primarily** to kill or harm humans.
 2. Humans, not robots, are responsible agents. Robots are tools designed to achieve human goals.
 3. Robots should be designed in ways that assure their safety and security.
 4. Robots are artifacts; they should not be designed to exploit vulnerable users by evoking an emotional response or dependency. It should always be possible to tell a robot from a human.
 5. It should always be possible to find out **who is legally responsible** for a robot.

¹ Koocher GP, Keith-Speigel P. Ethics in psychology and the mental health professions: standards and cases (Oxford textbooks in clinical psychology). USA: Oxford University Press; 2008.

SOURCE CODES OF CONDUCT

- The three laws of (fictional) robotics:
 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.



CHANGING LEGISLATION

REGULATING NEW DEVICES |

- In the US, the FDA recognizes medical devices and permits their sale.
 - About 99% of new devices are **cleared** if they are “**substantially equivalent**” to existing devices.
 - Otherwise, despite guidance released in 2012, new devices must go through very rigorous “**premarket approval**”, sometimes requiring clinical trials. Devices then fall into three classes:
 - **Class I** devices are **low risk**; they do not support or sustain life.
 - E.g., *dental floss*
 - **Class II** devices **do not cause harm** if used as intended.
 - E.g., *acupuncture needles, power wheelchairs*.
 - **Class III** devices are **high risk** and subject to the highest scrutiny
 - E.g., *replacement heart valves*.


REGULATING NEW DEVICES 2

- The software that powers data-collection devices was often proprietary, rather than open source (i.e., they weren't open to external scrutiny and auditing).
- While a recently-granted exemption to the Digital Millennium Copyright Act provides the opportunity to examine code for *external* medical devices, it may be even more important to examine *internal* medical devices, which are currently excluded from the exemption¹.

¹ Karen Sandler, Lysandra Ohrstrom, Laura Moy and Robert McVay, “Killed by Code: Software Transparency in Implantable Medical Devices,” Software Freedom Law Center, New York, 2010

STRATEGIES

- The [Affordable Care Act](#)'s shifts from a fee-for-service to a pay-for-performance model¹
 - Health IT is *rewarded*, at least in Medicare.
- Despite prohibitions in the Genetic Information Non-discrimination Act of 2008, there is growing interest in using genetic risk information for insurance stratification².
 - Differential pricing has become one of the standard practices for data analytics vendors, introducing new avenues to perpetuate inequality.
- The (current) White House views AI as providing “increased medical efficacy, patient comfort, and less waste”³.



¹ David Blumenthal, Melinda Abrams, and Rachel Nuzum, “The Affordable Care Act at 5 Years,” *New England Journal of Medicine* Vol. 372, Issue 25, (2015): 2453

² Yann Joly et al., “Life Insurance: Genomic Stratification and Risk Classification,” *European Journal of Human Genetics* 22 No. 5 (May 2014): 575–79).

³ Bryan Biegel, & Kurose, J. F. (2016). *The National Artificial Intelligence Research and Development Strategic Plan*.

H.R.6 – 114TH CONGRESS 21ST CENTURY CURES ACT |

- The [21st Century Cures Act](#) passed House of Representatives (344-77).
 - 13 July 2015: Received in the Senate, read twice, and referred to the Committee on Health, Education, Labor, and Pensions.
- **Guidance I**, “general wellness products”: Include “audio recordings, video games, software programs and other products that are commonly ... available from retail establishments.”
 - The FDA will not regulate such products as medical devices, as long as they meet two factors: they i) are intended for only general wellness; and ii) present low risk to users.
 - Such a device may claim that it “may help to reduce the risk of” or “may help living well with” certain chronic diseases. An acceptable claim for a software product might be that it “coaches breathing techniques and relaxation skills, which, as part of a healthy lifestyle, may help living well with migraine headaches.” The product’s value derives from *information*, rather than doing something directly to the body.

H.R.6 – 114TH CONGRESS 21ST CENTURY CURES ACT 2

- **Guidance II**, “real world evidence” (RWE): Although not usually used to win approval of a new device, RWE can be used to gain the FDA’s permission for a device to be used for more indications than the one for which it was originally approved.
 - What is the source of data to build the evidence? “The data is typically derived from electronic systems used in health care delivery, data contained within medical devices, and/or in tracking patient experience during care, including in home-use settings.”
- **Guidance III**, adaptive design of clinical trials supporting the FDA’s approval of new medical devices. “Adaptive” refers to “a clinical study design that allows for prospectively planned modifications based on accumulating study data without undermining the study’s integrity and validity.”
 - If *poorly* executed, adaptive design risks ‘moving the goalposts in the middle of the game’, posing hidden risks to patients. If *well* executed, adaptive design can reduce the time and cost of clinical research.
 - How is this data going to be collected? H.R.6 allows for more sharing of research data, possibly in response to HIPAA.

ADMINISTRIVIA

QUIZ AND PRESENTATIONS

- **Quiz 2** will be on Friday 25 November and take ~ 40 minutes (5% of overall mark).
- **Presentations** will take place Friday 9 December (10% of project mark).
 - 10 minutes presentation + 5 minutes question, per project (we start at 10h, sharp).
 - Asking questions will count towards your 'participation' mark.
 - Design, clarity, and structure (stay on time!) are of primary importance.
 - Outline i) goals, ii) state-of-the-art, iii) methodology, iv) any results or ongoing work.
 - Each member should speak
 - Please send your .pptx, .pdf, or .keynote by end of 8 December.

REPORTS

- Report is
 - 15% **data analytics**: *a good understanding/presentation of data; descriptive statistics*
 - 15% **code**: *you should submit everything you have. Try to keep it clean/correct.*
 - 30% **experiments** and analysis: *show meaningful comparisons; test your assumptions;*
 - 15% **literature review**: *show synthesis; contrast and compare; recency and relevancy.*
 - 15% **technical quality**: *motivation and understanding*
 - 10% **overall presentation**: *spelling, structure, figures/tables*