Enabling Trust, Accountability, and Routine Use of AI-Enabled Healthcare

Richard Giordano, Reham Al Tamime, Peter West
Web Science Institute

The HEALTH-I Team

UNIVERSITY OF Southampton

NEXT++

Challenges facing healthcare

We are living longer! But, this means more chronic illness.

Diabetes
422 million worldwide
Almost 4x more than 1980

(Mathers 2006)

Heart failure
6.5 million in USA
Predicted to rise 46% by 2030

(American Heart Association 2017)

Doctors are facing **increasing workload** and a need for more **personalised care**.

Typical clinical scenarios

Example scenario:

- Patient presents to a doctor with heart palpitations (irregular heart rate)
- She has had palpitations for the past few weeks
- Doctor examines the patient and asks about their medications and lifestyle
- Doctor then refers patient for an EKG
- Three weeks latest, EKG doesn’t reveal anything.
- Doctor concludes the palpitations were caused by atrial fibrillation.

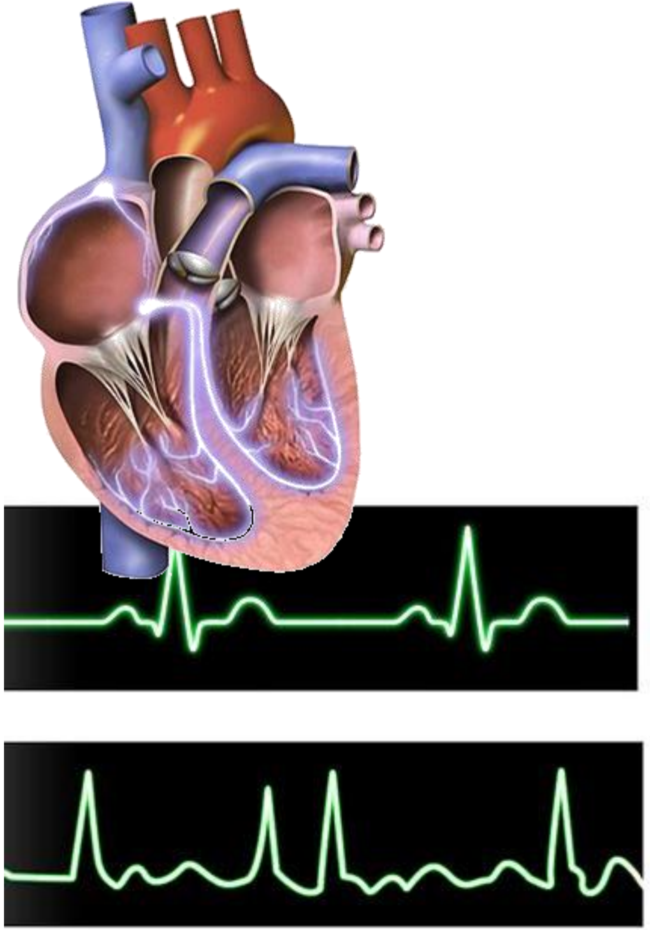

Where could AI help medicine?

AI could help doctors get a better understanding of patients

For our patient, AI could help the doctor identify the cause of the patient’s palpitations.

But where can we get the data?

Patient-Generated Data

Any kind of data which a patient has recorded using their own means.

Wearables
Fitbit, Apple Watch

Smartphone apps
Google Fit, Strava

Health products
Blood pressure cuffs, weighing scales

Journals
Hand-written and electronic

Health Self-Tracking Tools are Increasingly Popular

One third of US adults track at least one indicator of health (such as weight or symptoms) on using an app (MobiHealth News 2013)

Over **15 million Fitbits** sold in first quarter 2017 (Statista 2018)

Photo by Phillip Pessar

Food tracking

Apps like DietLens help people track their calorie intake by analysing photos of food.

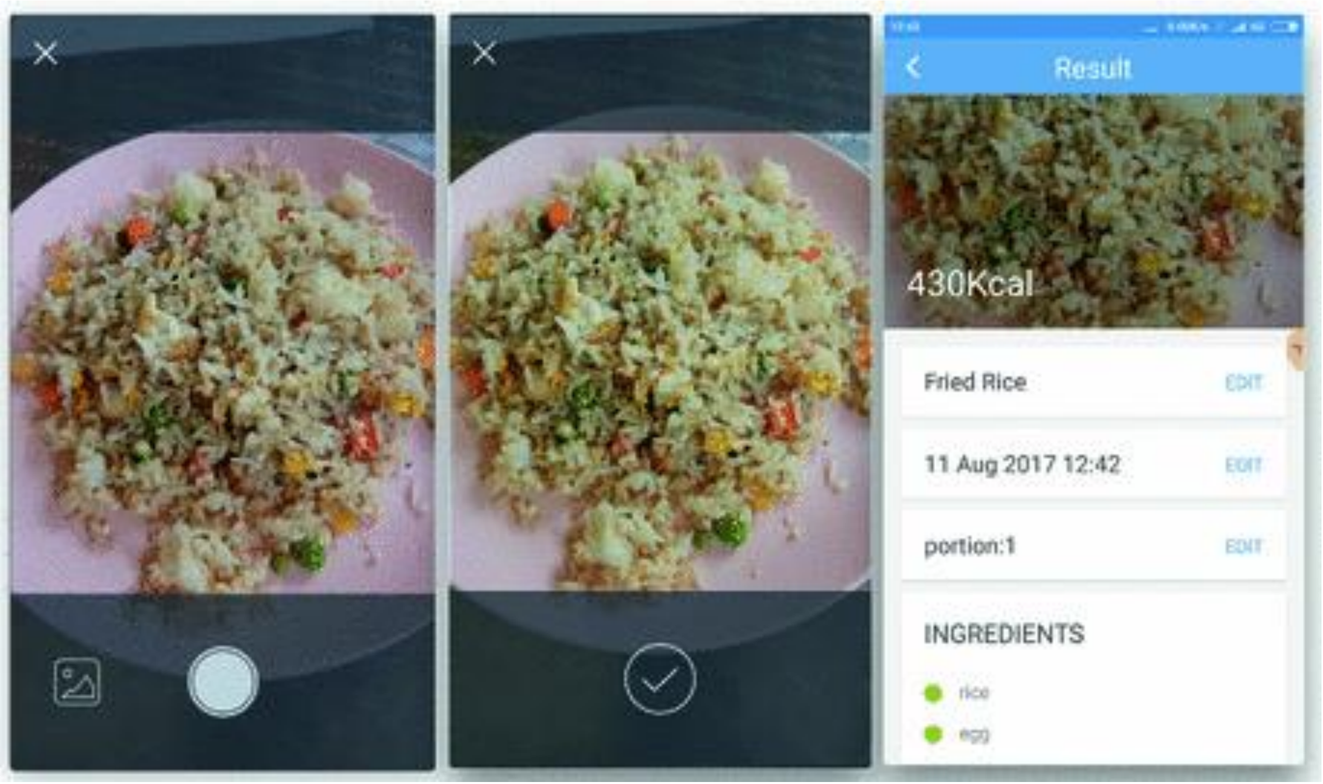

Ming ZY., Chen J., Cao Y., Forde C., Ngo CW., Chua T.S. (2018) Food Photo Recognition for Dietary Tracking: System and Experiment. In: Schoeffmann K. et al. (eds) MultiMedia Modeling. MMM 2018. Lecture Notes in Computer Science, vol 10705. Springer, Cham

We asked 13 clinicians about the future of healthcare

Clinical role	Participants	Years in practice
Cardiologist	P1, P2, P3, P4	All 20+ years
Mental health specialist	P5, P6	10 years, 5 years
Emergency doctor	P7	5 years
Junior surgeon	P8	5 years
Hospital doctor	P9	4 years
General practitioner	P10	20+ years
Heart failure nurse	P11	20+ years
Oncology nurse	P12	2 years
Audiologist	P13	3 years

All were practicing in the UK

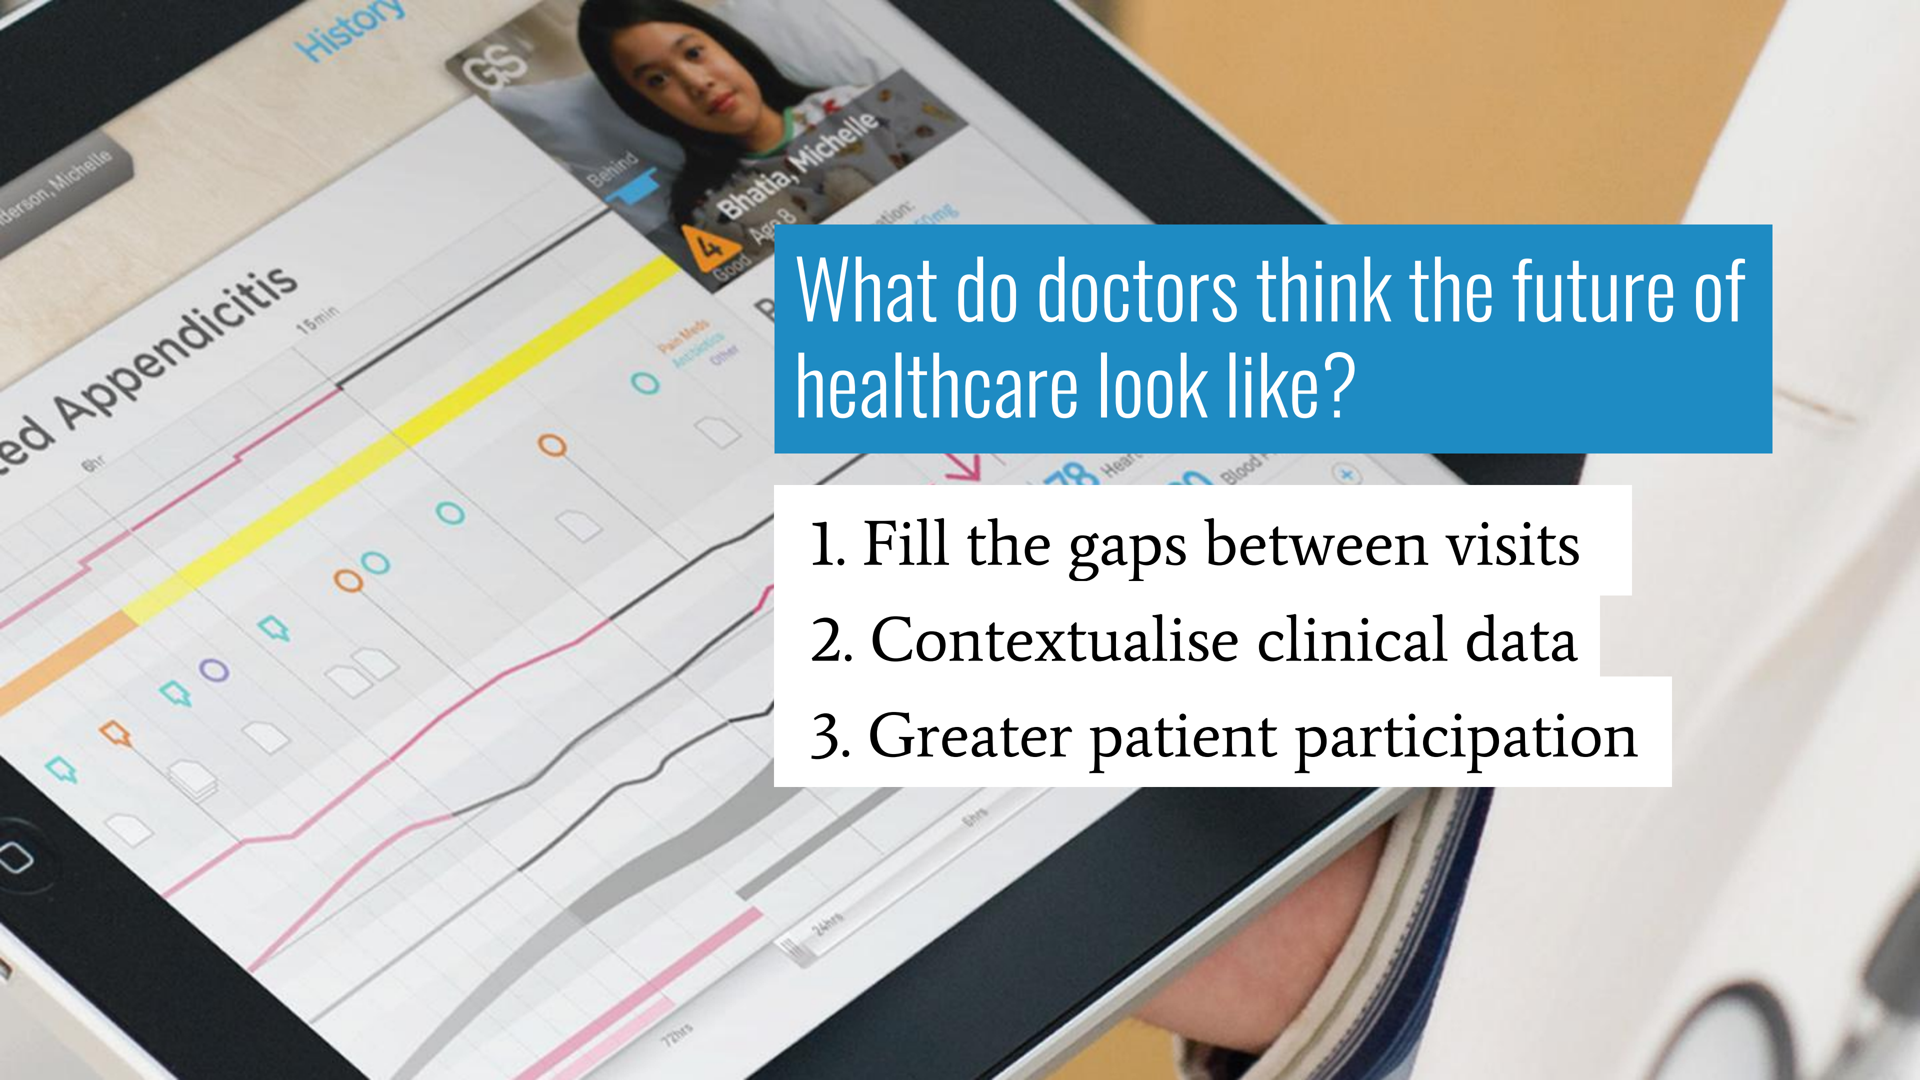

What do doctors think the future of healthcare look like?

1. Fill the gaps between visits
2. Contextualise clinical data
3. Greater patient participation

In the hospital of the 2050?

Demo: http://flamingtempura.github.io/pgd-view

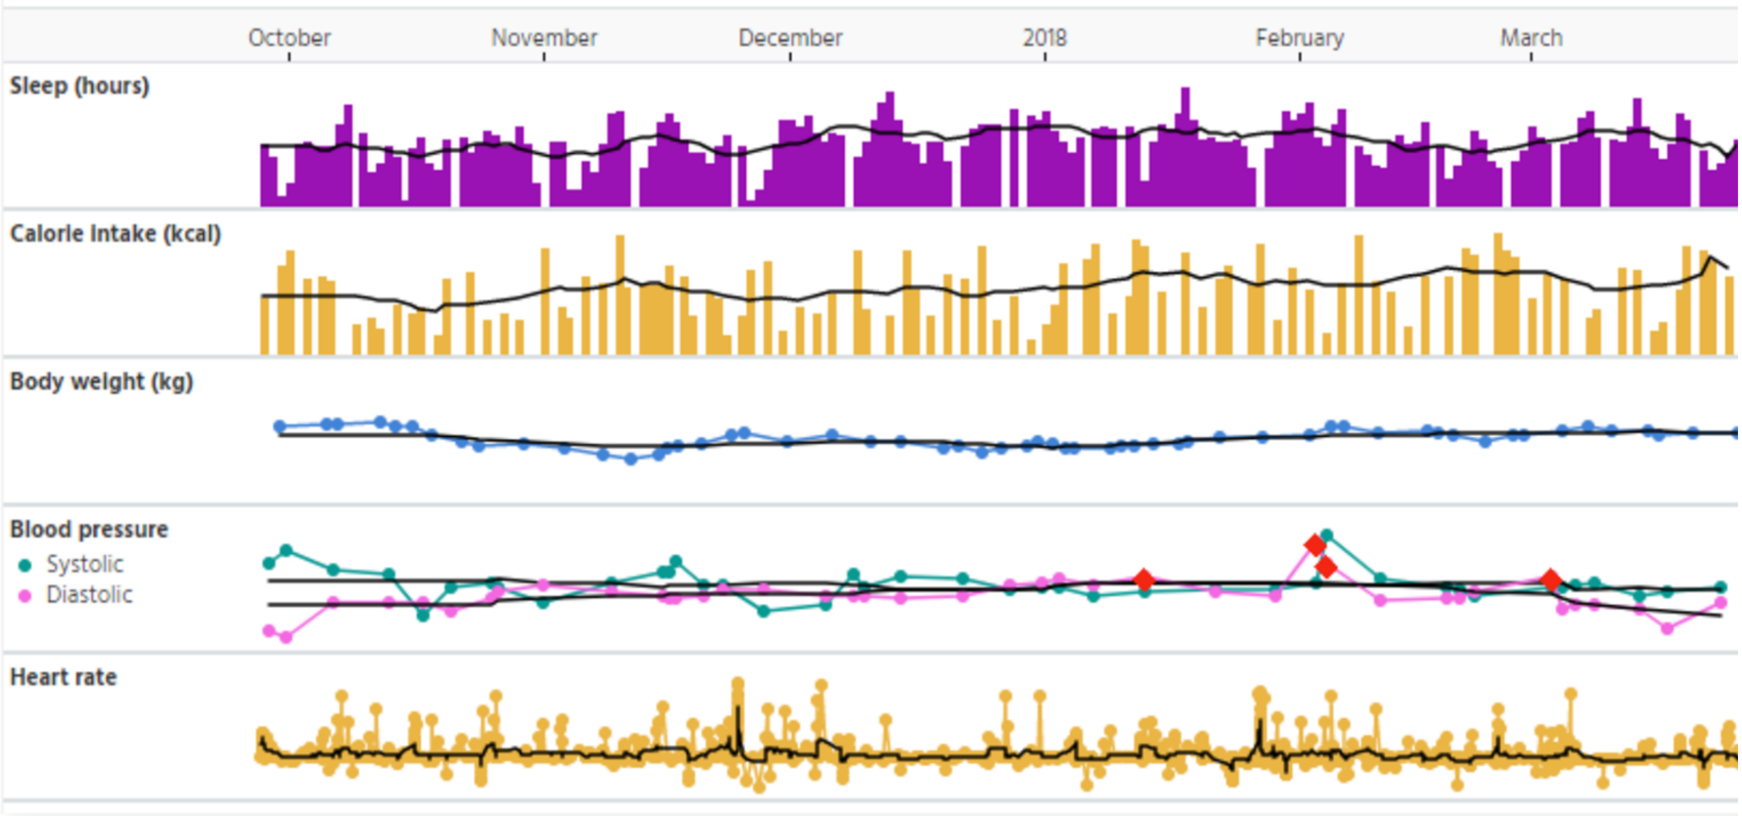

In the hospital of the 2050?

Monday 10am: palpitations

- 2 hours sleep

Tuesday 12pm: palpitations

- 4 hours sleep

Thursday 2pm: palpitations

- 3 hours sleep

Poor sleep was leading to worse palpitations.

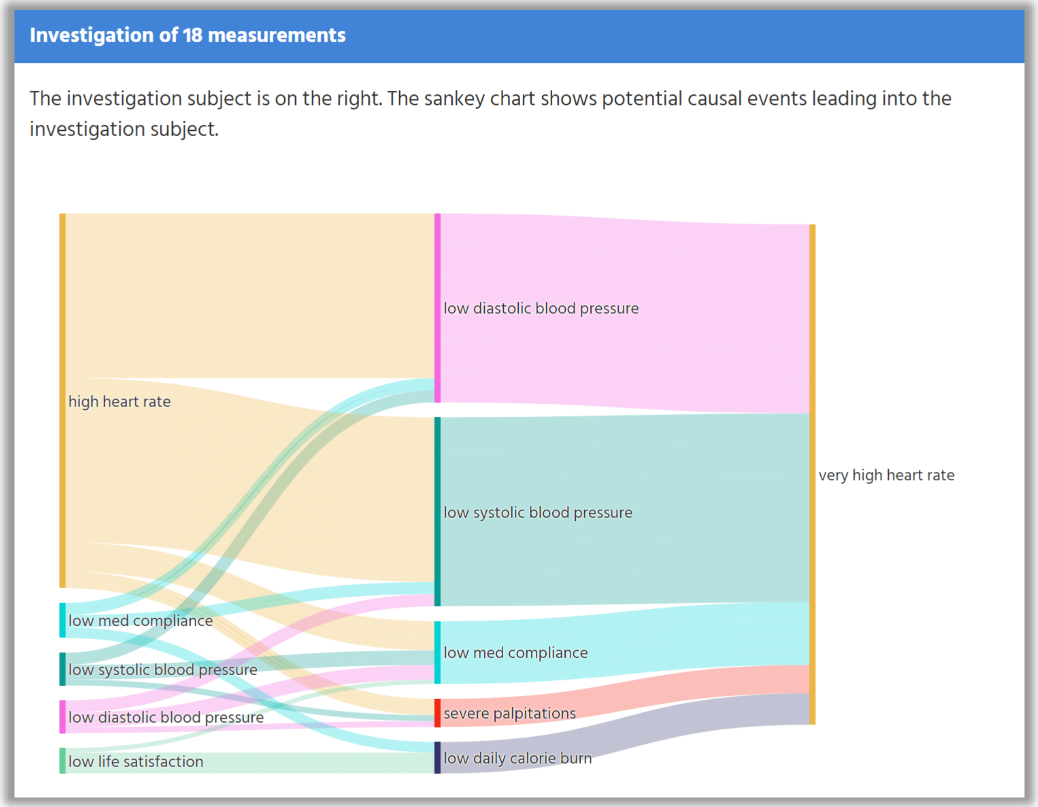

But how will doctors perceive AI?

Some doctors have resisted the idea of using AI in healthcare.

Will patient privacy be upheld?

Will patients trust it?

WIRED

WIRED Security

AI has no place in the NHS if patient privacy isn't assured

DeepMind is working on a technical solution to boost transparency when it comes to AI in healthcare – but it's a long road to machines gaining patient trust

By **NICOLE KOBIE**

Understanding the privacy concerns of mobile health apps users

Reham Al Tamime

Research Problem

- The privacy decisions become more complex.
- Lack of understanding of users' privacy concerns about the data collection and sharing practices in different contexts.
- Users are still provided with traditional privacy management options.

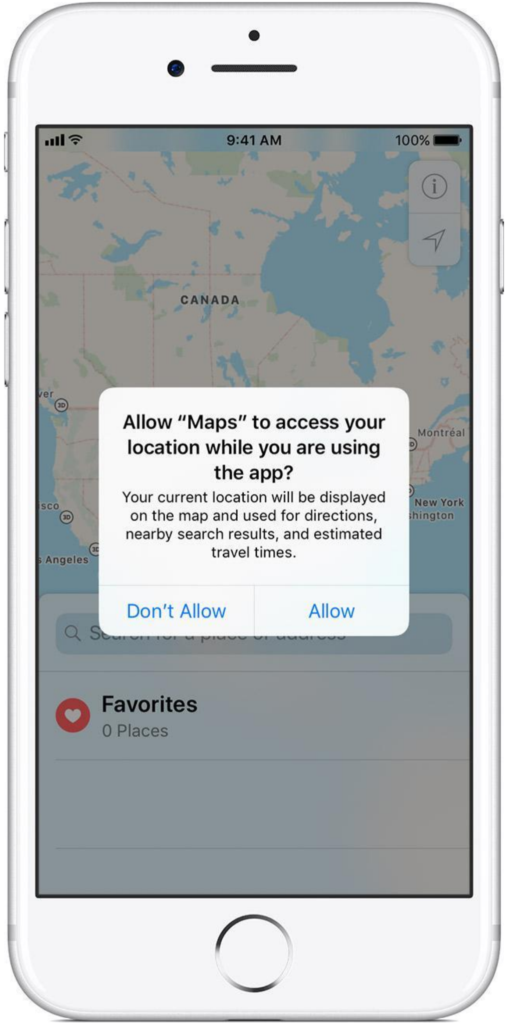

Research Goals

- To understand the privacy concerns of mobile health app users:

  a. To understand the comfort level in sharing data across various contexts.
  b. To investigate how crowdsourcing can help to understand privacy preferences.

Theory of contextual integrity

- Contextual integrity ties adequate protection for privacy to norms of specific contexts [Helen Nissenbaum, 2004].
- Information gathering and dissemination be appropriate to that context and obey the governing norms of distribution within it.
- Capturing and specifying context elements of data collection and sharing are of great importance.

Scenarios of contexts

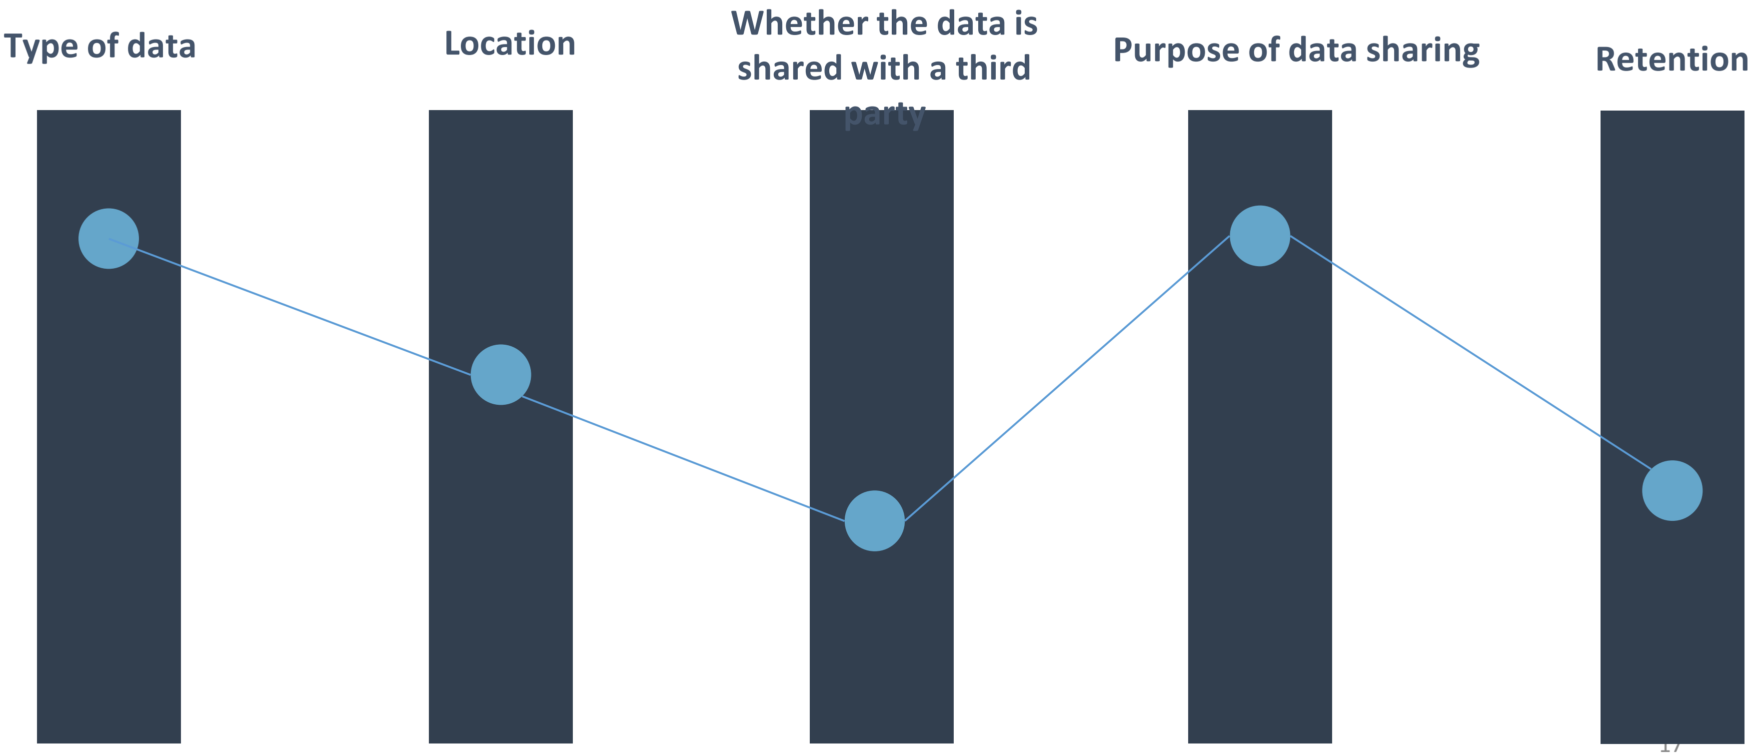

Zooniverse

- Crowdsourcing, by definition, is the use of humans (at scale) to complete computationally difficult or time consuming tasks.

- Crowdsourcing for citizen science has been used to help annotate scientific datasets, such as Hubble Telescope images.

- 432 scenarios in total -> 16 subsets -> 27 scenario each.

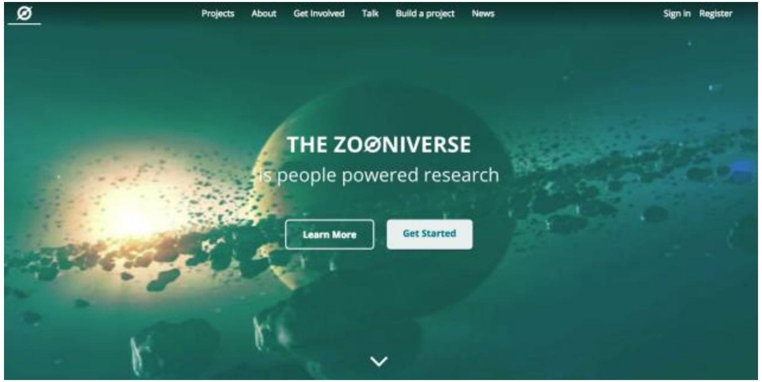

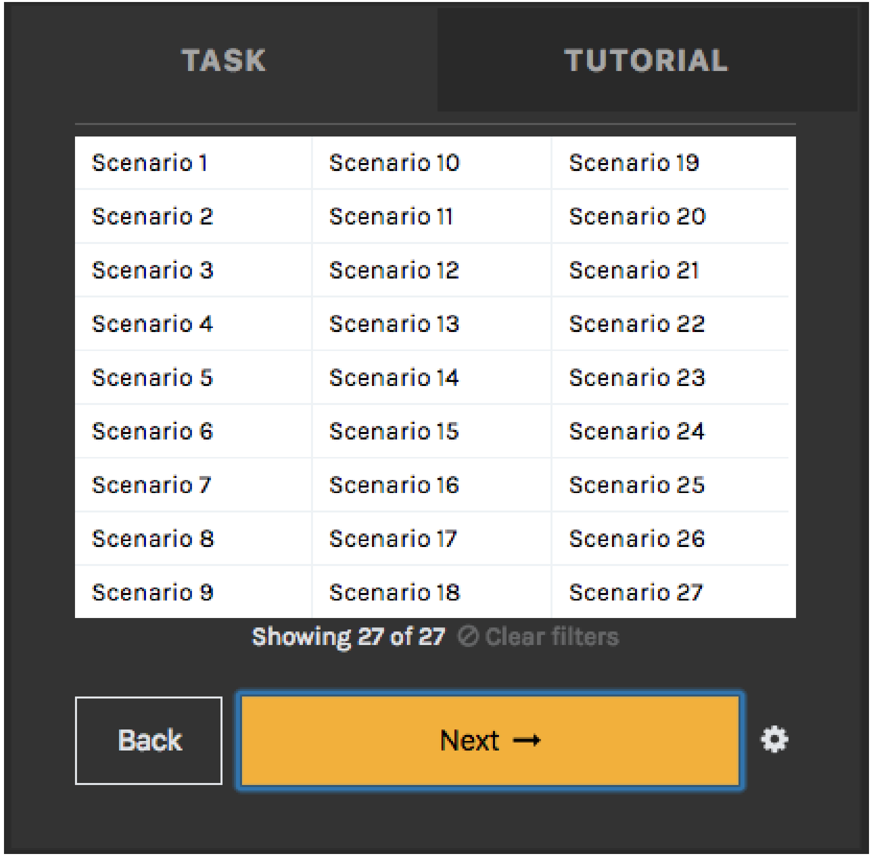
TASK
TUTORIAL
Scenario 1
Scenario 10
Scenario 19
Scenario 2
Scenario 11
Scenario 20
Scenario 3
Scenario 12
Scenario 21
Scenario 4
Scenario 13
Scenario 22
Scenario 5
Scenario 14
Scenario 23
Scenario 6
Scenario 15
Scenario 24
Scenario 7
Scenario 16
Scenario 25
Scenario 8
Scenario 17
Scenario 26
Scenario 9
Scenario 18
Scenario 27
Showing 27 of 27
Clear filters
Back
Next →

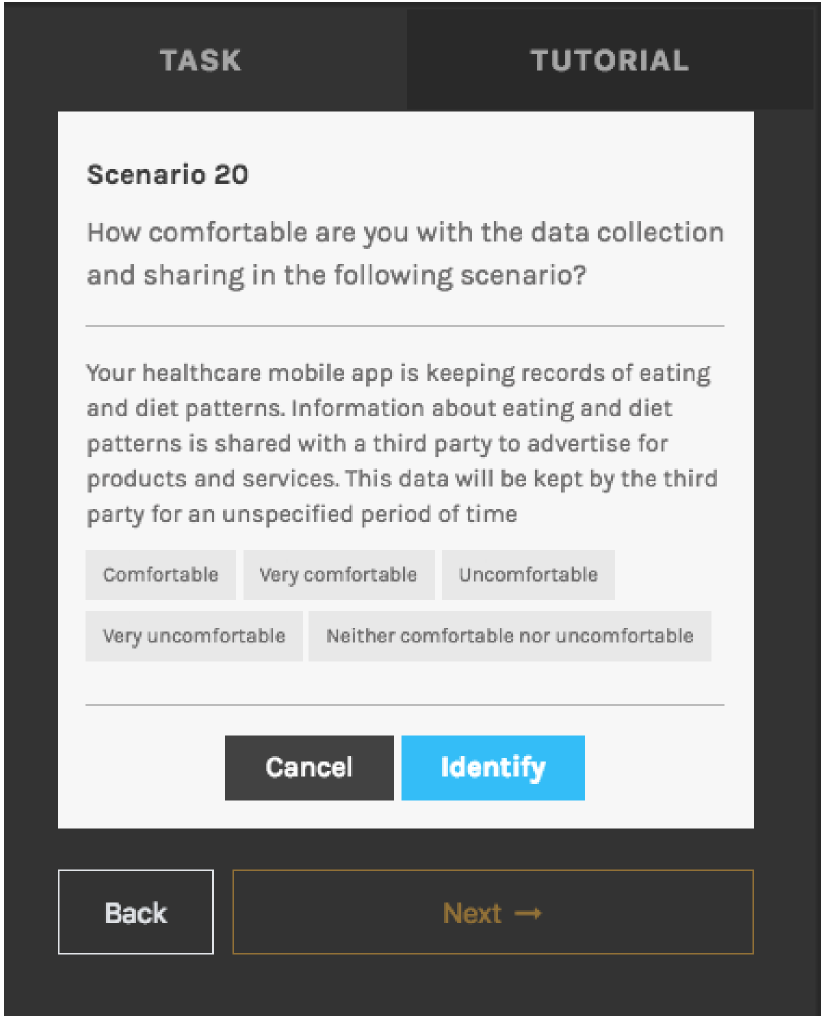
TASK
TUTORIAL
Scenario 20
How comfortable are you with the data collection and sharing in the following scenario?
Your healthcare mobile app is keeping records of eating and diet patterns. Information about eating and diet patterns is shared with a third party to advertise for products and services. This data will be kept by the third party for an unspecified period of time
Comfortable
Very comfortable
Uncomfortable
Very uncomfortable
Neither comfortable nor uncomfortable
Cancel
Identify
Back
Next →

Trust in AI healthcare technology

- Build framework that provides crowd-based privacy support for sharing data in various contexts.

- Design AI technologies that take into consideration users' privacy concerns.

Black Box Medicine

Richard Giordano

Personalized Medicine

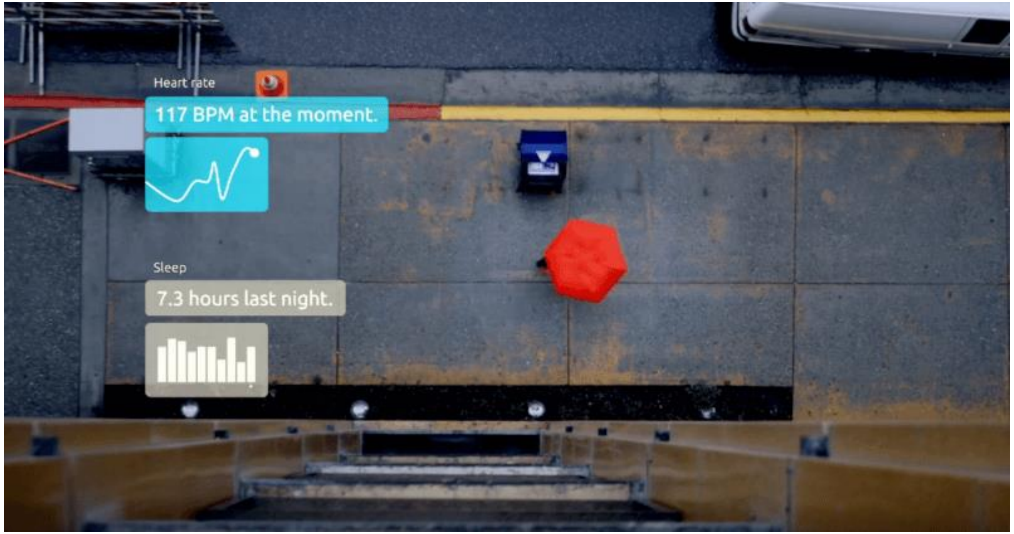

One-Size-Fits-All works only partially

Fail to respond to treatment

38% depression

40% asthma

75% cancer

Use of massive datasets and machine learning to design treatments for individuals

algorithms + data structures = programs

- Niklaus Wirth (1976)
  - Algorithms and data structures are intimately related
  - If you want to sort, use arrays
- Opacity of algorithms
  - Trade secrets
  - Technical literacy
  - Characteristics of the algorithm and the scale required to apply them usefully
  - National Nurses Union
    - “Algorithms are simple mathematical formulas that nobody understands”

Algorithmic programming: (PL/I)

```
                SOURCE LISTING

STMT LEV NT

   1       0  BLKLUF: PROC OPTIONS (MAIN) REORDER;
              /* SELECTS ALL PATRON CHARGES WITH INTERNAL STATUS OF */
              /* 7, 8, OR 9 WHICH ARE DUE BEFORE MARCH, 1985.       */
              /* LOOKS UP ID'S IN LUF.  DOES NOT PROCESS RECORDS    */
              /* WHERE LUF PRIMARY STATUS IS "X".  REARRANGES       */
              /* NAME SO THAT LAST-NAME IS FIRST.                   */
   2   1   0  DCL MARIMST FILE RECORD INPUT;
   3   1   0  DCL OUTFILE FILE RECORD OUTPUT;
   4   1   0  DCL SYSMSG EXTERNAL PRINT FILE;
   5   1   0  DCL SYSPRINT EXTERNAL PRINT FILE;
   6   1   0  DCL PATRONS KEYED RECORD INPUT FILE ENV (VSAM);
   7   1   0  DCL RESPONSE CHAR (9) INIT ('         ');
   8   1   0  DCL FNDIT BIT (1);
   9   1   0  DCL PAGE FIXED BIN INIT (1);
  10   1   0  DCL 1 MARIREC ALIGNED,
                  2 L1 FIXED BINARY (15),
                  2 FLAGS BIT (8),
                  2 TC_TM BIT (8),
                  2 ST2_NCT BIT (8),
                  2 OLDSEQ CHAR (3),
                  2 CHG FIXED BINARY (20),
                  2 DUE FIXED BINARY (20),
                  2 STAT1 CHAR (1),
                  2 ID CHAR (9),
                  2 BOOK CHAR (8),
                  2 LOC CHAR (2),
                  2 EC$A CHAR (48);
  11   1   0  DCL 1 IDREC BASED (PTR),
                  2 ID CHAR (9),
                  2 SOURCE CHAR (1),
                  2 STAT1 CHAR (1) INIT (' '),
                  2 STAT2 CHAR (1),
                  2 L3 FIXED BIN,
                  2 LINES (10 REFER (L3)) CHAR (28);
  12   1   0  DCL 1 OUTREC,
                  2 LNF CHAR (28),
                  2 X0 CHAR (1) INIT (' '),
                  2 L_STAT1 CHAR (1) INIT (' '),
                  2 X1 CHAR (1) INIT (' '),
                  2 M_STAT1 CHAR (1) INIT (' '),
                  2 X2 CHAR (1) INIT (' '),
                  2 M_ID CHAR (9),
                  2 X3 CHAR (1) INIT (' '),
                  2 BFLG CHAR (1) INIT (' '),
                  2 X4 CHAR (1) INIT (' '),
                  2 BOOKNMBR CHAR (8),
                  2 X5 CHAR (1) INIT (' '),
                  2 EC$A CHAR (48),
                  2 X6 CHAR (1) INIT (' '),
                  2 DUE PICTURE '999999';
  13   1   0  DCL MARIRECS FIXED DECIMAL (6) INIT (0);
```

```
PL/I OPTIMIZING COMPILER        BLKLUF: PROC OPTIONS (MAIN) REORDER;

  STMT LEV NT

   14   1  0  DCL I, J FIXED DEC (6) INIT (0);
   15   1  0  DCL (MOREMARIRECS) BIT (1) INIT ('1'B);
   16   1  0  DCL TEMP CHAR (28) VAR INIT (' ');
   17   1  0  DCL LNF CHAR (28) INIT (' ');
   18   1  0  ON KEY (PATRONS) FNDIT = '0'B;
   19   1  0  ON ENDFILE (MARIMST) MOREMARIRECS = '0'B;
   20   1  0    FNDIT = '0'B;
   21   1  0  READ FILE (MARIMST) INTO (MARIREC);
   22   1  0    RESPONSE = '         ';
   23   1  0   TEMP = ' ';
   24   1  0  LNF = ' ';
   25   1  0  DO MARIRECS = 1 BY 1 WHILE (MOREMARIRECS);
   26   1  1   IF MARIREC.LOC = 'OJ' THEN MOREMARIRECS = '0'B;
   27   1  1   IF MOD (MARIRECS, 1000) = 0 THEN PUT SKIP FILE (SYSMSG)
                 EDIT ('->PROCESSING RECORD: ', MARIRECS)
                 (A, P'ZZZ,ZZ9');
   28   1  1    RESPONSE = MARIREC.ID;
   29   1  1    FNDIT = '1'B;
   30   1  1  OUTREC.BFLG = ' ';
   31   1  1  IF (MARIREC.DUE < 850228) THEN DO;
   32   1  2        IF ((SUBSTR(MARIREC.ST2_NCT, 1, 4) = '1001'B) |
                        (SUBSTR(MARIREC.ST2_NCT, 1, 4) = '1000') |
                        (SUBSTR(MARIREC.ST2_NCT, 1, 4) = '0111'B)) THEN
               IF VERIFY (SUBSTR(MARIREC.ID, 6, 3), '0123456789') = 0 THEN DO;
   33   1  3    FNDIT = '1'B;
   34   1  3    READ FILE (PATRONS) SET (PTR) KEY (RESPONSE);
   35   1  3  IF IDREC.STAT1 ¬= 'X' THEN DO;
   36   1  4    IF FNDIT THEN DO;
   37   1  5     DO I = 28 BY -1 TO 1
                           WHILE (SUBSTR(IDREC.LINES(1), I, 1) = ' ');
   38   1  6      END;
   39   1  5     TEMP = SUBSTR(IDREC.LINES(1), 1, I);
   40   1  5     DO J = I-1 BY -1 TO 1
                        WHILE (SUBSTR(TEMP, J, 1) ¬= ' ');
   41   1  6      END;
   42   1  5     LNF = ' ';
   43   1  5        LNF = SUBSTR(IDREC.LINES(1), J) ||
                          SUBSTR(IDREC.LINES(1), 1, J);
   44   1  5    END;
   45   1  4  IF ¬FNDIT THEN DO;
   46   1  5      LNF = ' ';
   47   1  5      IDREC.STAT1 = ' ';
   48   1  5      END;
   49   1  4          OUTREC.LNF = LNF;
   50   1  4          OUTREC.L_STAT1 = IDREC.STAT1;
   51   1  4          OUTREC.M_STAT1 = MARIREC.STAT1;
   52   1  4          OUTREC.M_ID = MARIREC.ID;
   53   1  4          OUTREC.BFLG = SUBSTR(MARIREC.FLAGS, 4, 1);
   54   1  4          OUTREC.BOOKNMBR = MARIREC.BOOK;
   55   1  4          OUTREC.EC$A = MARIREC.EC$A;
   56   1  4          OUTREC.DUE = MARIREC.DUE;
   57   1  4          WRITE FILE (OUTFILE) FROM (OUTREC);
   58   1  4       END;
   59   1  3    END;
```

```
PL/I OPTIMIZING COMPILER        BLKLUF: PROC OPTIONS (MAIN) REORDER;

  STMT LEV NT

   60   1  2   END;
   61   1  1   MARIREC.EC$A, OUTREC.EC$A = ' ';
   62   1  1   READ FILE (MARIMST) INTO (MARIREC);
   63   1  1  END; /* ENDS THE MAJOR LOOP */
   64   1  0  PUT SKIP FILE (SYSMSG) EDIT (
                   '*> E N D  O F  P R O C E S S I N G ') (A);
   65   1  0  RETURN;
   66   1  0  END BLKLUF;
```

The Invisible Maniac was algorithmic

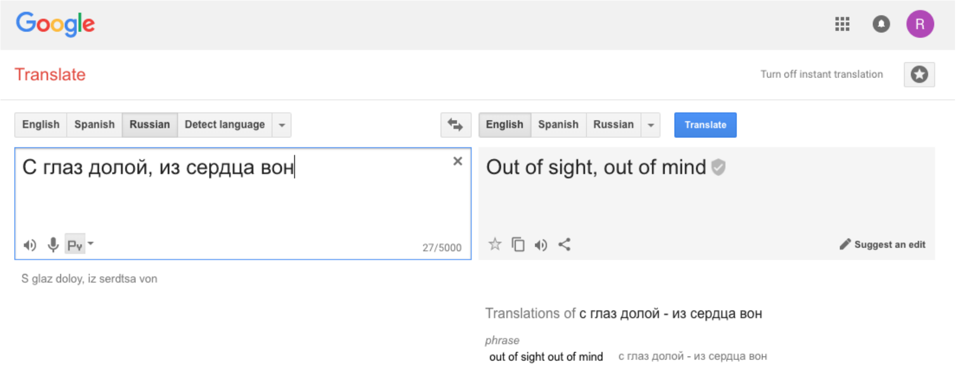

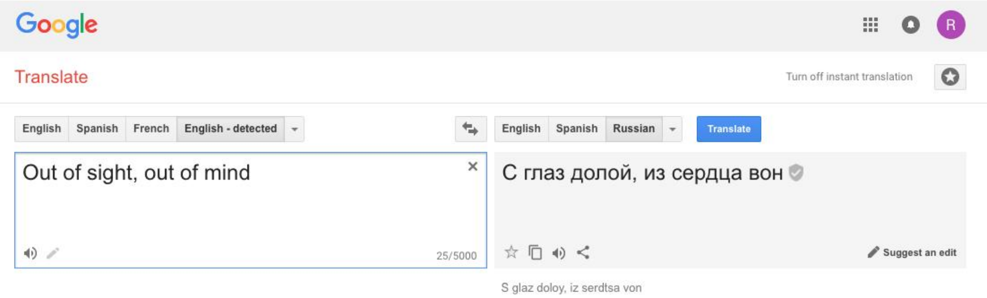

Machine learning

- Machines trained from data
    - Classifiers produce categories
    - Learners train on data (based on models) and produce weights
    - Inductive reasoning

- “Applications that cannot be programmed by hand”

Black Box Medicine

- Opaque computational models to make decisions or judgements related to health care
- Use of large scale datasets and associated algorithms (and heuristics) to use implicit, complex connections between multiple patient characteristics
- Algorithms are neither explicit nor transparent
- Relationships they capture cannot be explicitly understood
- Relationships often cannot be explicitly stated
- This is not deliberately hidden—it is a consequence of complexity

Challenges for the Patient

- How to explain treatment plan to a patient
  - Patients often do not understand information or retain it
- Subgroups have different levels of trust in healthcare
  - White women | African American women
  - Native born | Immigrants
- Predictive categorization not based on who/what you are
  - A viewer likely to enjoy a movie (Netflix)
  - A customer likely to buy this item (amazon)
  - A teenager likely to commit a crime (NYC predictive policing)
  - A women likely to become pregnant
  - A genotype likely to respond to CBT to treat schizophrenia

Challenges for the Doctor

- Treatments and care pathways rely on complex biological interactions through integration of of information from interdisciplinary fields holistically
  - Common in Systems Biology
  - Not part of mainstream clinical research
    - How do we develop/train doctors?
- Quality of machine learning relies on aspects of training data and models
  - Who is responsible?
- Deductive reasoning in medicine
  - Test a theory empirically-–randomized clinical trials
- Inductive reasoning (Black Box Medicine)
  - Pattern recognition
  - Inductive reasoning not trusted among medics since it yields false positives

Challenges

- Policy
  - Equity across populations
- Institutions
  - Governance structures
- Technical
  - Systems to produce an audit trail
    - (This is not trivial...)

Conclusion: Enabling Trust, Accountability, and Routine Use of AI-Enabled Healthcare

Policy and societal challenges relating to privacy, trust, and transparency

(We can't just throw programmers at the problem)

Our challenge to the NExT++ Workshop:

How can we build the next generation of *accountable* AI?

Our Published Research

Al Tamime, Giordano R, Hall, W (2018) Observing Burstiness in Wikipedia Articles during New Disease Outbreaks **Web Science Conference, Amsterdam, Netherlands**

West P, Van Kleek M, Giordano R, and Weal M (2018) Common barriers to the use of patient-generated data across clinical settings. **CHI 2018, Montreal, Canada**

West P, Van Kleek M, Giordano R, Weal M, and Shadbolt N (2017) Information quality challenges of patient-generated data in clinical practice. **Frontiers Public Health**.

West P, Giordano R, Van Kleek M, and Shadbolt N (2016) The quantified patient in the doctor's office: Challenges and opportunities. **CHI 2016, San Jose, USA. (Honorable Mention)**

See our posters outside!