“Recent Developments and Future Possibilities in Precision Health”

Michael R. Kosorok

W.R. Kenan, Jr. Distinguished Professor and Chair

Department of Biostatistics

Professor

Department of Statistics and Operations Research

University of North Carolina at Chapel Hill

Recent Developments and Future Possibilities In Precision Health

Michael R. Kosorok, PhD

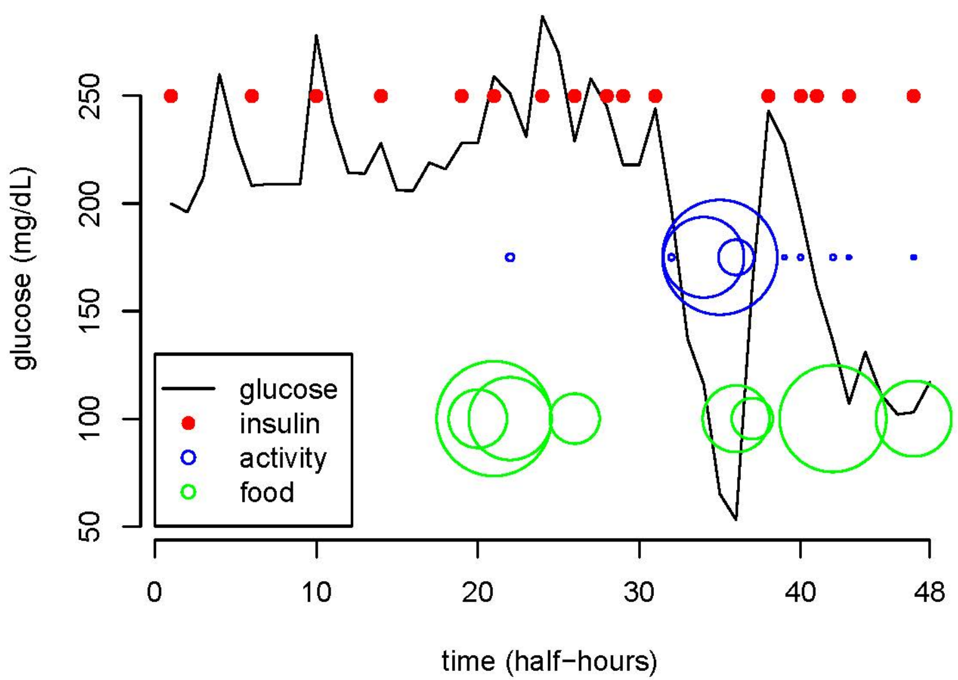

Time course plot of glucose, Insulin, physical activity, and food intake for a single patient (data from Maahs et al, 2012)

May 1, 2020

UNC GILLINGS SCHOOL OF GLOBAL PUBLIC HEALTH

Precision Medicine and Precision Health

- Precision Medicine
  - Data-driven decision support for treating patients
  - Treatment can include drug choice, administrative actions, dosing, timing, potentially modifiable risk factors, etc.
  - Must be reproducible and generalizable
- Precision Public Health
  - Data-driven decision support for families, clinics, communities, hospitals, social networks, etc.
  - Potential treatments can also be policies
- Precision Health
  - Union of precision medicine and precision health
  - We always consider the consequences of our actions on the populations involved not just narrow subgroups or individuals

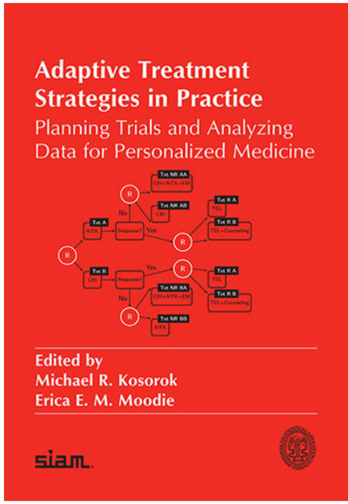

Operating Principles

- Observable Constituents:
  - Tailoring variables (X)
  - Choice of treatments and/or potentially modifiable risk factors (A)
  - Vector of outcomes or utilities (Y)
  - Could be multiple (X,A,Y) triples over time for each patient
- Dynamic Treatment Regime (DTR):
  - Single decision: make a single recommendation for treatment
  - Multiple decision: make a series of interdependent recommendations
  - Continual monitoring: for diabetes, mHealth
- Role of Heterogeneity in the data:
  - Heterogeneity of patients is beneficial (essential) for good precision medicine analysis so that our treatment rules are broadly applicable
  - Need heterogeneity of treatment assignment (either naturally or by design) in the data so we can determine best treatment under a variety of situations

Overall Pipeline Details

- Dynamic Treatment Regime:
  - d(X) gives recommended A to maximize Y in future patients
  - Regression: model Y as a function of X and A (Q(X,A)=E[Y|X,A] is the "value"), with interaction between X and A being most important
  - Policy estimation: directly estimate d(X) without Q(X,A) (e.g., outcome weighted learning)
  - Prediction versus prescriptive decision support:
    - Suppose Y=f(X)+Ag(X)+e, where bigger Y is better and A={0 or 1}
    - We only care about g(x), since rule d(X)={1 if g(X)>0, 0 otherwise} yields optimal
    - A focus on prediction yields too much focus on f(X) instead of g(X)
- Propensity Score:
  - P(A|X) is propensity score
  - We can estimate P(A|X) from data and make sure it is positive for all X
- Causal Methods:
  - Potential outcome validity (Y(0) and Y(1) are well behaved)
  - No unmeasured confounders: Y(0) and Y(1) are independent given X
  - Positivity assumption (P(A|X)>0 for all X)

DTR Estimation

- Single-decision setting:
  - Regression and maximization
  - Frequentist and Bayesian tools can be used (both good, but some non-Bayesian is better when data is high dimensional)
  - Machine learning is very helpful
- Multi-decision setting:
  - Need off-policy reinforcement learning
  - Q-learning is especially useful
  - Q-learning involves a sequence of standard regressions: random forests are good here
  - There is software for this (in R mostly)

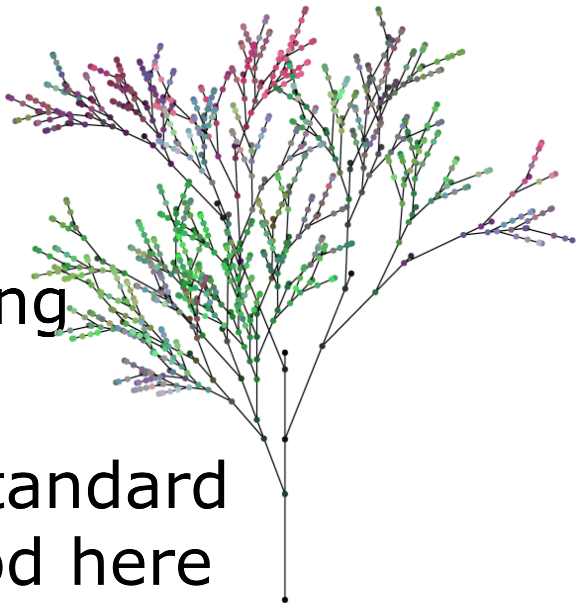

Some Recent Developments

- Outcome weighted learning (OWL):
  - Focus on estimating policy rather than entire process
  - OWL: convert to classification problem of predicting A from X weighted by Y (avoids regression model)
- Benefits of OWL:
  - Robust to regression model specification
  - Converts problem to a classification task which can employ machine learning (such as SVM)
  - Performs very well in many situations
  - Many extensions

What is Artificial Intelligence Relative to ML?

What is intelligence?

- To think or act like people?
  - The Turing test (1950)
  - Spelling errors
  - Taking time to solve math
- To perform tasks people can perform?
- To reason logically?
- To be able to make on optimal decision?
- Learning how to act optimally?
- To reason morally?
- To create something surprising?

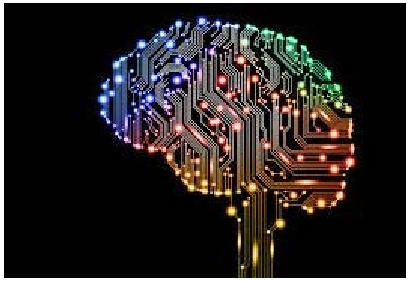

What do we want computers to do?

- Make rational decision to achieve pre-defined goals
- Goals are usually related to optimizing a predefined utility
- Can a computer decide what utility it should use?

The three stages of prediction

- Prediction
- Causally valid prediction enabling “what if” analysis
- Decision support: what is the best action to take in this context?

Beware of Winter

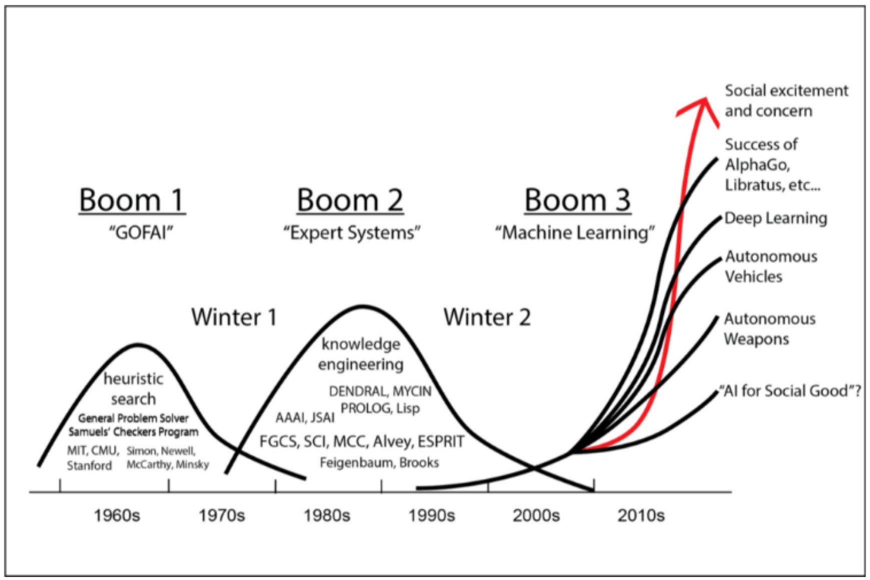

The dark side of AI history

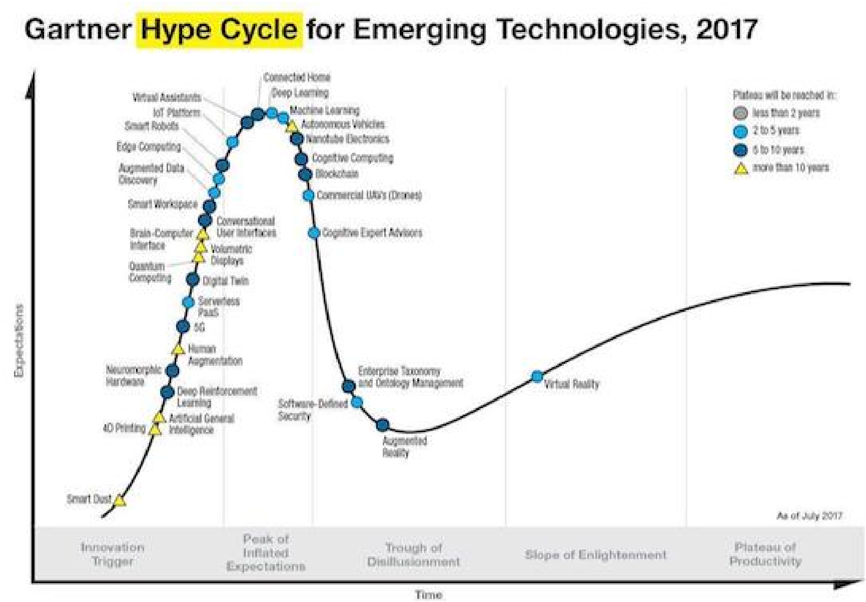

Typically, AI advances in stages. Winters can be mitigated if we stay realistic and scientifically responsible.

The Dark Side of AI/ML

The democratization of AI

- Easy to use software is widely available
- Pros: this can lead to greater awareness of AI and citizen science progress
- Cons: insufficient understanding of statistical principles (reproducibility), too many false positives

Good research takes care

- Easy implementation leads to impatience in research
- This leads to inadequate evaluation (peer review)
- Promising advances may be dismissed early due to poor implementation
- Study design is crucial

The risk of bad actors

- Evil people could use AI for harm
- Automated AI tools could have unintended adverse consequences:
  - Unintended disparities
  - Provoking discord/conspiracy theories

Example 1: T1D and Safe Exercise

- Goal: precision medicine for minute-by-minute control of blood glucose level and overall control of weight gain
- Framework: data-driven decision science at multiple time scales using new hybrid SMART designs, mHealth, micro-randomization, system dynamics, and statistical control theory (robotics)

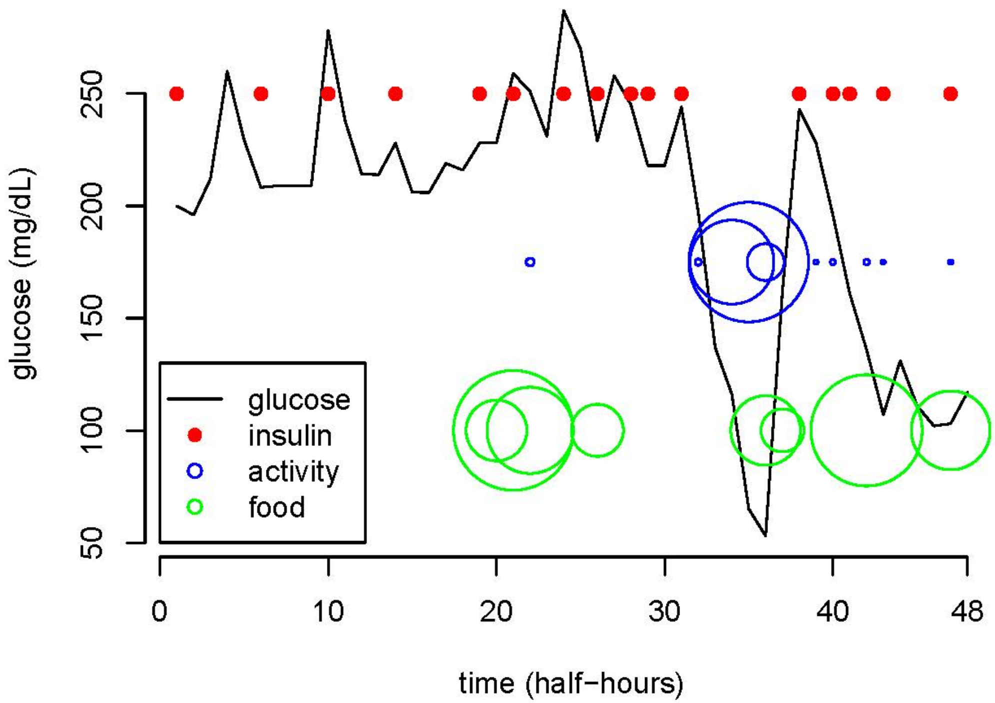

Time course plot of glucose, insulin, physical activity, and food intake for a single pediatric patient (data from Maahs et al, 2012)

Example I, cont.

- We use a Markov decision process mathematical and statistical framework
- We measure the state of the patient and actions which have been taken
- V-learning is a new kind of reinforcement learning which incorporates outcome weighted learning

We assume the data consists of a sample of *n* i.i.d. trajectories $(S^1, A^1, S^2, \ldots, S^T, A^T, S^{T+1})$ where *S* is state, *A* is action (treatment), and *U* is a specified utility (outcome)

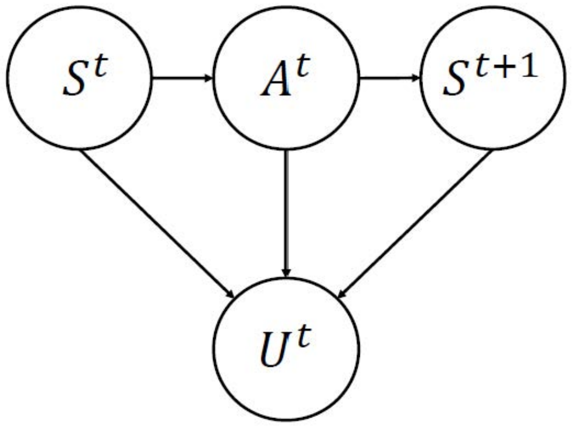

Graphical depiction of a Markov decision process indexed by discrete time points

Results From V-Learning

- We applied V-learning to the mobile health data of Maahs et al (2018)
- State included insulin use, food intake, physical activity, and blood glucose
- Possible actions included take insulin, eat, exercise, combinations (such as take insulin and eat)
- Outcome: weighted average time glucose out of range
- Simulations verify that V-learning leads to significantly better performance than Q-learning (≈ twice as good)
- For this data, the DTR estimated from V-learning reduced the time out of range by 64%
- *Good optimization tools are crucial*

Example 2: Bipolar Disorder

- The Systematic Treatment Enhancement Program for Bipolar Disorder Standard Care Pathway (STEP-BDE SCP)
- Challenging to treat:
    - Characterized by episodes of depression and mania
    - Anti-depressants can treat depressive episodes
    - However, anti-depressants may induce mania episodes
- Clinical decision making needs to balance trade-offs between depression and mania

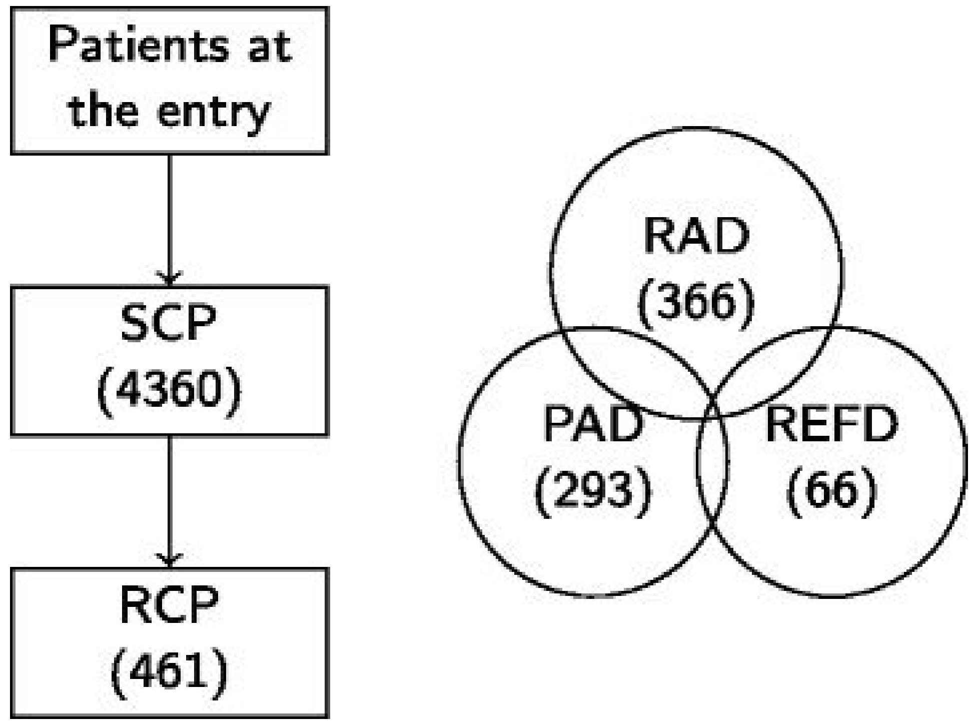

STEP-BD Design. Patients in SCP satisfying RCP criteria entered RCP. In RCP, there are three different pathways: RAD, PAD, and REFD. (Wu, Laber, Lipkovich and Severus, 2015). 1437 of SCP patients had bipolar disorder.

Example 2, cont.

- We used a modification of inverse reinforcement learning.
- The idea is that we assume that the physicians properly balance the two outcomes (depression and mania) with a non-zero probability.

$\sqrt{n}$ convergence but non-standard limiting distribution requiring empirical processes and the Argmax theorem, with unusual non-standard bootstrap needed for inference

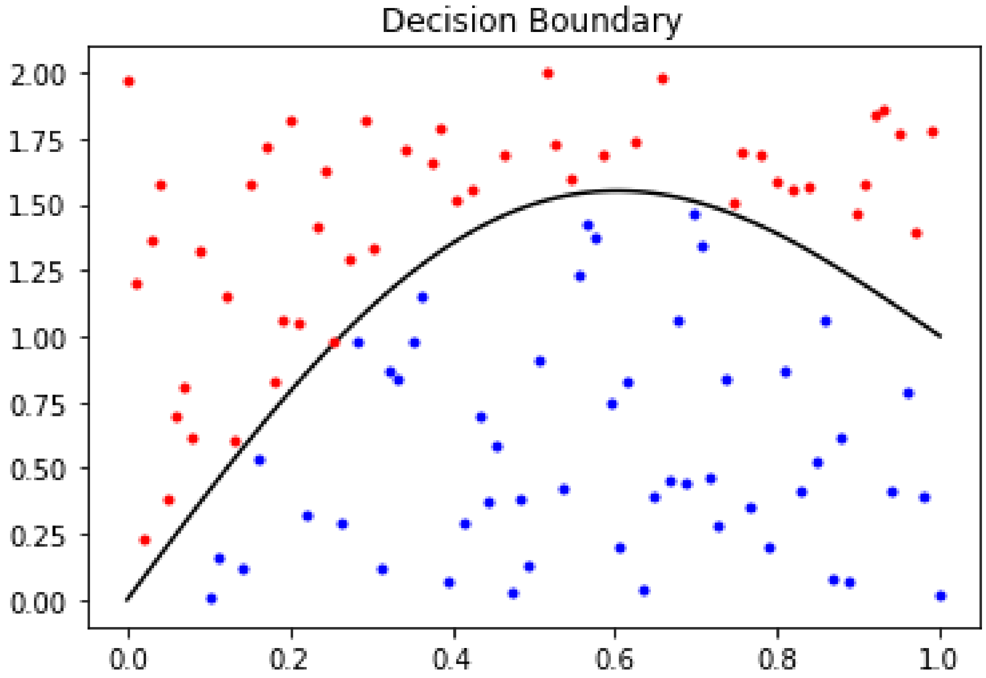

(Rome, 2017). Not actual decision boundary.

UNC GILLINGS SCHOOL OF GLOBAL PUBLIC HEALTH

Results From STEP-BDE SCP Analysis

- Analysis of the observational study with 1437 patients having bipolar disorder (Sachs et al, 2007, *NEJM*).
- Using our proposed method, we were able to estimate an improved decision rule which led to a 7% improvement (p-value < 0.0001).
- Both increased age and history of substance abuse were important factors leading to lower recommended use of antidepressants.
- If we selected the two outcomes to be depression and side effect burden, we obtain an improvement of 9% (p-value < 0.001).

Looking to the Future

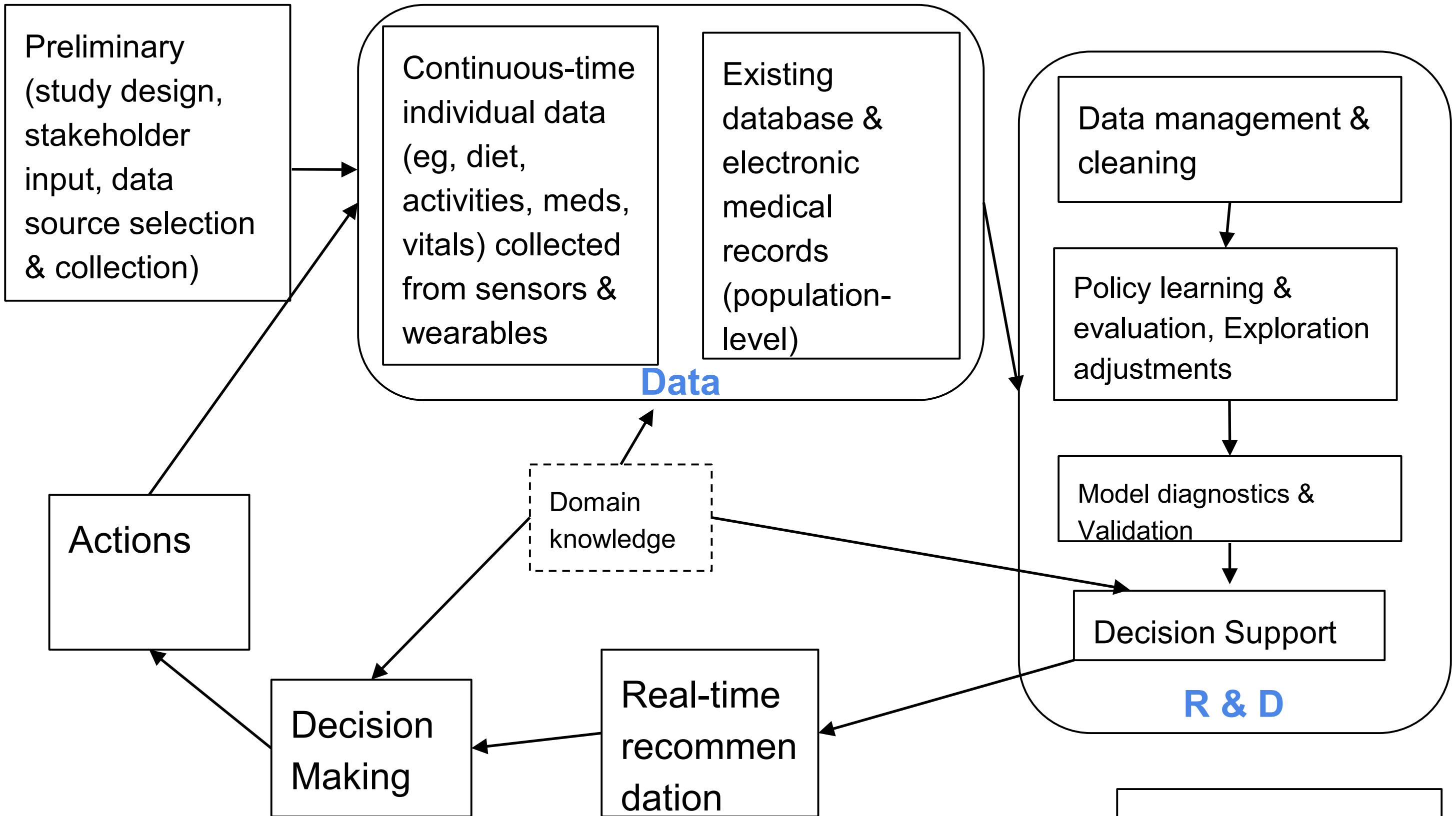

Decision Making With Continual Reassessment (Feedback Loop and Exploration)

Statistical Tools & Methodology Development (e.g. Machine Learning, Deep Learning, Reinforcement Learning, Causal Inference, Thompson sampling, etc.)

Computational Tools (e.g. software development and engineering, cloud computing, high-performance computing clusters, servers, GPUs, Apps, etc.)

Preparing for the Future

The best way to predict the future is to create it (Abraham Lincoln)

Do good AI/ML science

- Statistical foundation essential
  - Reproducibility
  - Generalizability
  - Study design
- Computing foundation essential
- Domain expertise essential
- Rigorous team science with depth

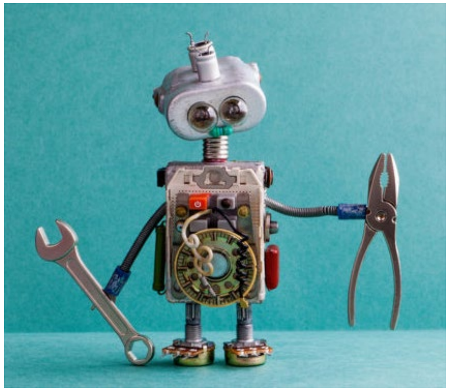

Training the next generation

- With increased democratization, teach fundamental statistical concepts
- All domain areas need knowledge of AI and data science
- Team science training

AI/ML ethics, etc.

- Ethics of AI are highly non-trivial
- Need to preserve privacy
- Need to break barriers to data use
- Need to study unintended consequences
- Need to develop strategies (e.g., game theory) to anticipate and reduce threats from bad actors

Communication about AI/ML

- Data science and AI experts need specialized communication training
- Everyone needs training in data science communication
- A new communications subdiscipline is needed which incorporates AI and data science
- Role of public opinion is crucial

Concluding Thoughts

- There are many useful new methods in AI and machine learning for precision health discovery
- We need to move away from prediction toward decision support
- Statistical inference principals are key to reproducibility and generalizability
- We need to embrace transdisciplinarity 
- Much work remains
- *We have not come this far just to have come this far*

Thank You