COLLEGE OF ENGINEERING
INDUSTRIAL & OPERATIONS ENGINEERING
UNIVERSITY OF MICHIGAN

Optimization in Medicine

INFORMS Healthcare Conference, Rotterdam, 2017

Brian Denton
Department of Industrial and Operations Engineering
University of Michigan

Optimization in Medicine

Cancer

Diabetes

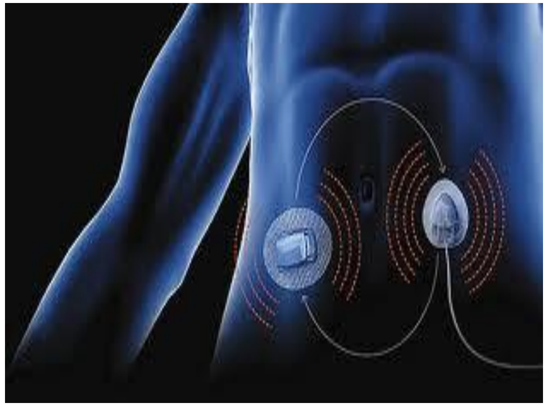

Kidney Disease

Heart Disease

Publications on PubMed in the last 10 years

- 1,426,842 articles on cancer
  - 169,076 articles on breast cancer
  - 79,662 articles on prostate cancer
- 474, 417 articles on heart disease and stroke
- 304,406 articles on diabetes
- 43,887 articles on kidney disease
- 2,935 articles on allergies

PubMed Results Over the Last 10 Years

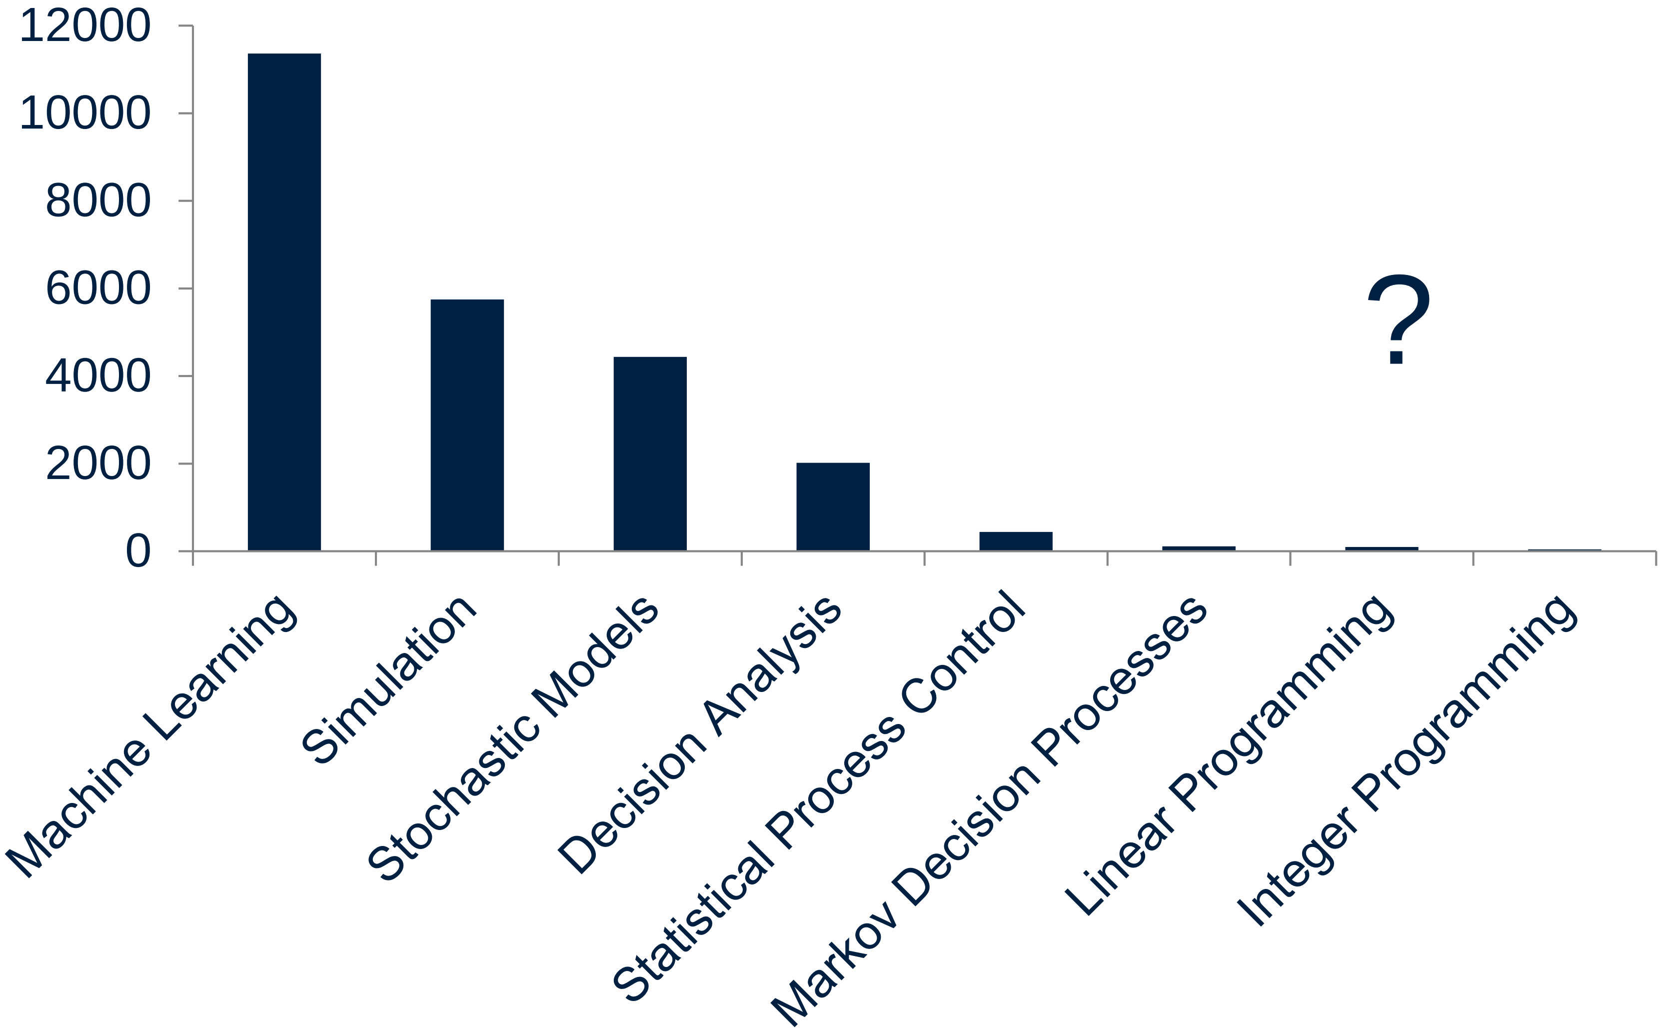

Ideas

1. Optimization can improve medical decision making

2. Medicine can improve optimization

3. There are many unaddressed opportunities for future impact

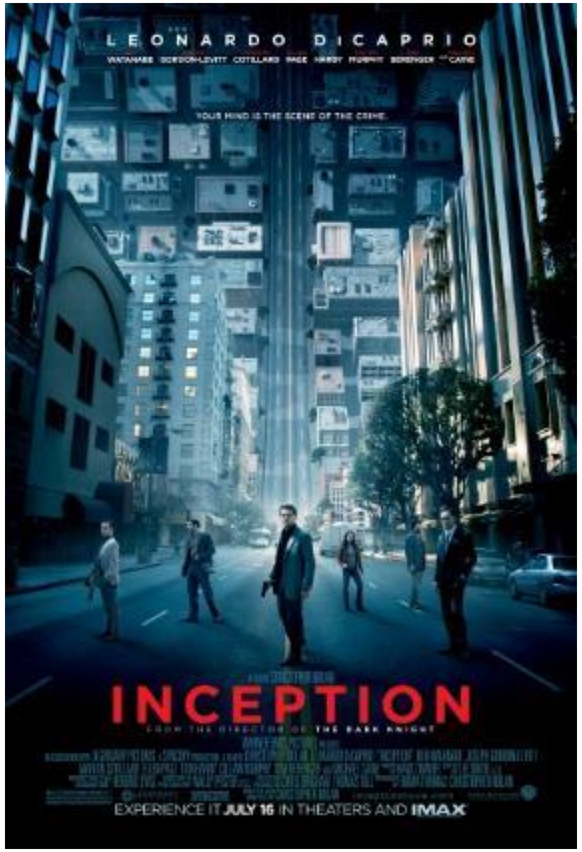


PRESENT
FUTURE
PAST

Example 1: Radiation Treatment

- External beam radiation is passed through the body harming cancerous and healthy tissue
- Objective: minimize damage to healthy tissue while delivering required dose to cancer tissue

Bahr et al, 1968, The Method of Linear Programming Applied to Radiation Treatment Planning, *Radiology*, 91, 686-693

Radiation is delivered via a rotating gantry with a multi-leaf collimator

Rotating Gantry

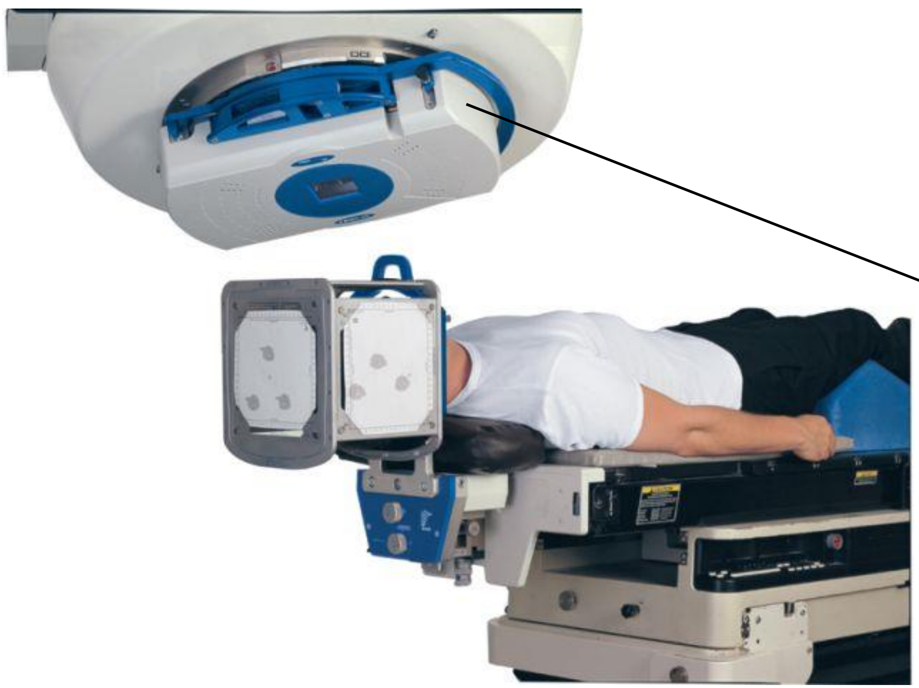

Multileaf Collimator

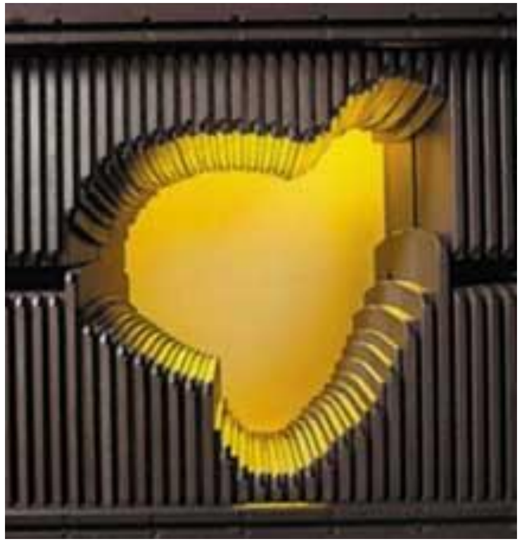

2-Beam Problem

1. Tumor
2. Spine
3. Brain

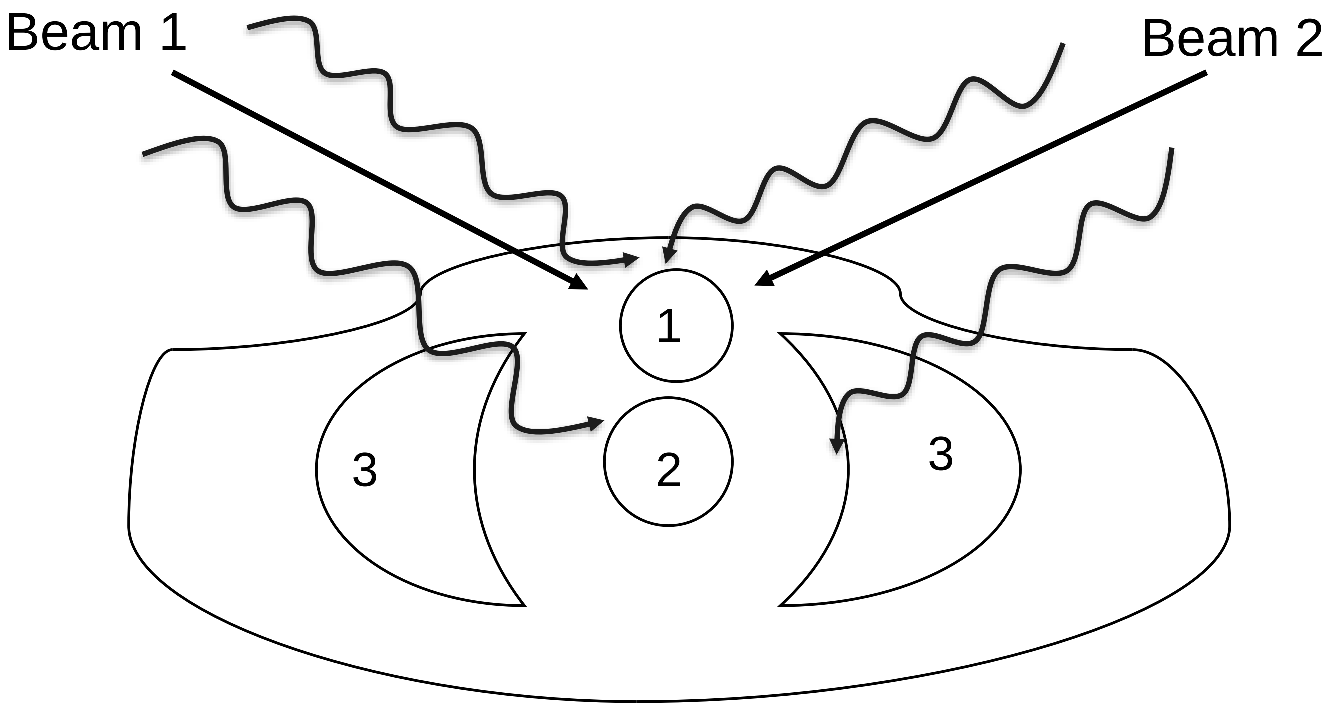

Adapted from “Optimization in Operations Research,” Ron Rardin

Linear Program

Decision Variables: Exposure times for beams 1 and 2 (x_1, x_2)

Area	Dose Absorbed		Restriction on Dosage in Kilorads
	Beam 1 Dose	Beam 2 Dose	
Brain	0.4	0.5	Minimize
Spine	0.3	0.1	$\leq$ 2.7
Tumor	0.5	0.5	$=$ 6
Center of tumor	0.6	0.4	$\geq$ 6

Linear Program

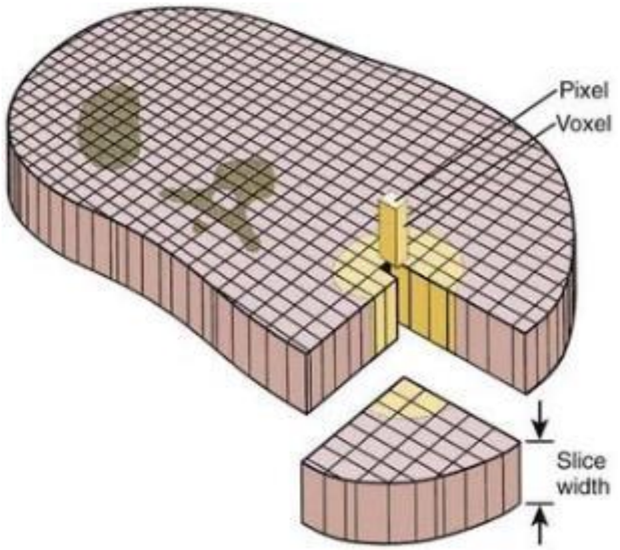

$$Min \ \ \sum_{\ell \in L} G_\ell(z)$$

$$Subject\ to:$$

$$z_j = \sum_{k \in K} D_{kj} x_k, \qquad for\ all\ j\ in\ V$$

$$x_k \geq 0, \ \ k \in K, \ \ z_j \geq 0, \ j \in V$$

z_j: the dose delivered to voxel $j \in V$

x_k: the duration of beam $k \in K$

Romeijn & Dempsey. (2008). Intensity modulated radiation therapy treatment optimization. *TOP*, *16*(2), 215.

Extensions

Integrated optimization of aperture design and beam intensities:

- Predefined number of beams

- Each beam is decomposed into a rectangular grid with m rows and n columns to create an intensity matrix

  - For each row there are $\frac{1}{2}n(n-1)+1$ combinations of left and right leaf settings

$$\Rightarrow \left(\frac{1}{2}n(n-1)+1\right)^m \quad \text{apertures}$$

- Column generation method: Start with a restricted set of apertures, price out new apertures (columns) via decomposition algorithms

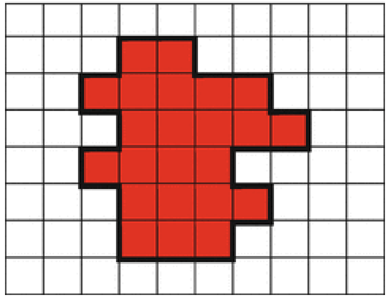

Romeijn, E., Ahuja, R., Dempsey, J.F., Kumar, A. 2005. "A column generation approach to radiation therapy treatment planning using aperture modulation." *SIAM Journal on Optimization* 15(3); 838-862.

Path from Research to Implementation

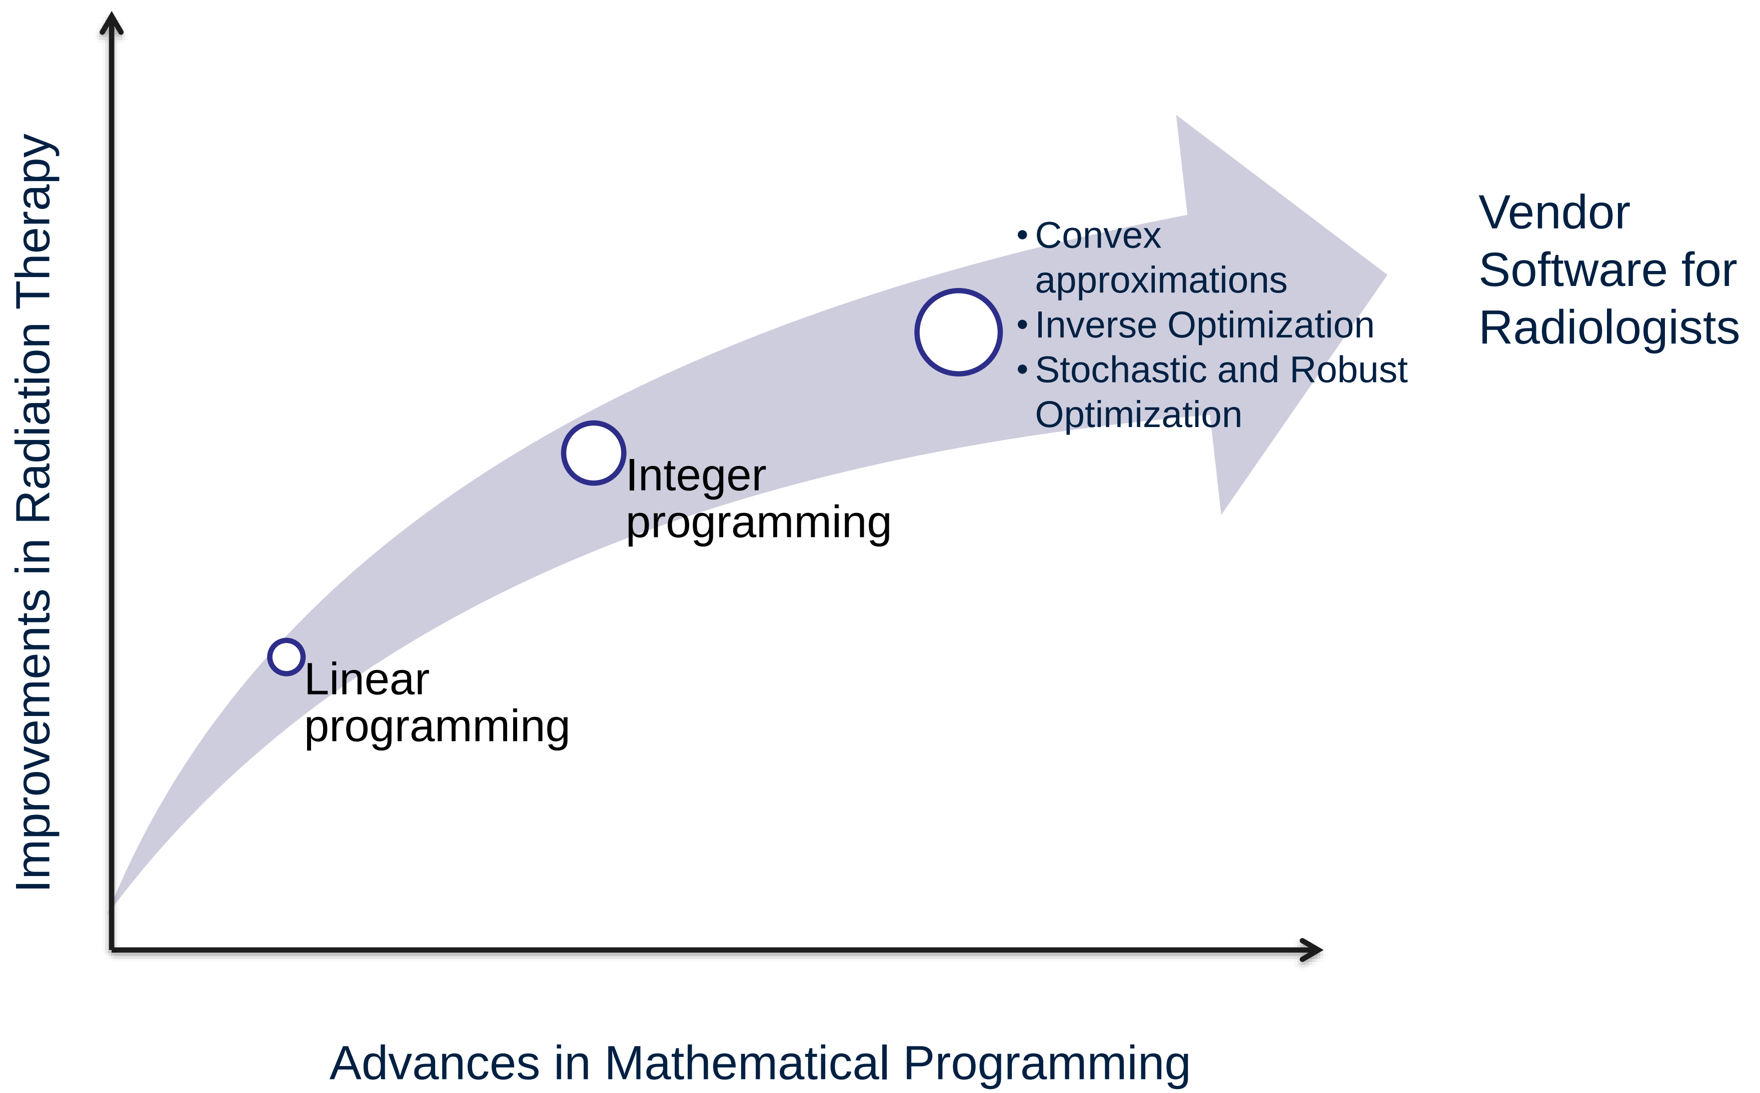

Example 2: Kidney Disease

- Principal treatment options:
  - Dialysis (home or clinic)
  - Transplant (live or deceased donor)

- More than 350,000 people are on dialysis and 80,000 waiting for transplant

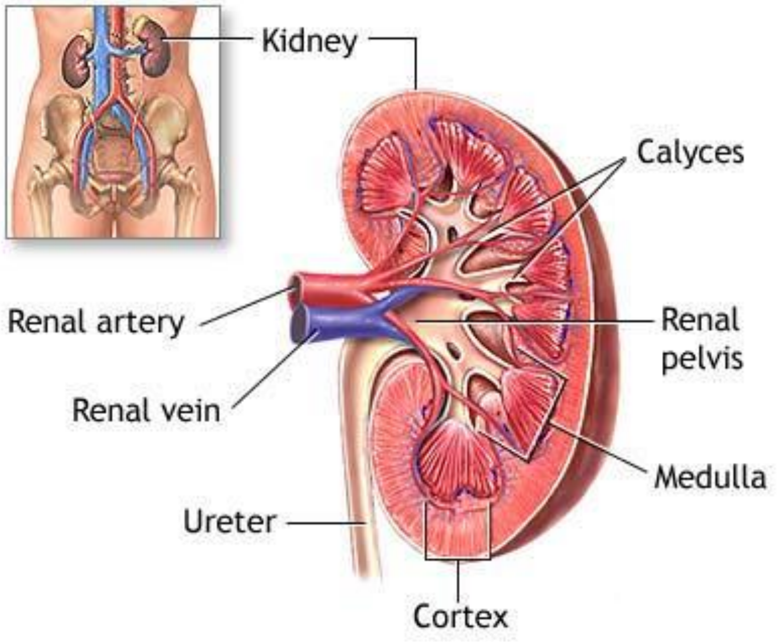

Nonlinear Optimization

Miller, J.H. et al. 1960. “Optimization of Certain Parameters in Hemodialysis,” *Transactions - American Society for Artificial Internal Organs*, 6(1); 68-75

Optimal time to change bath

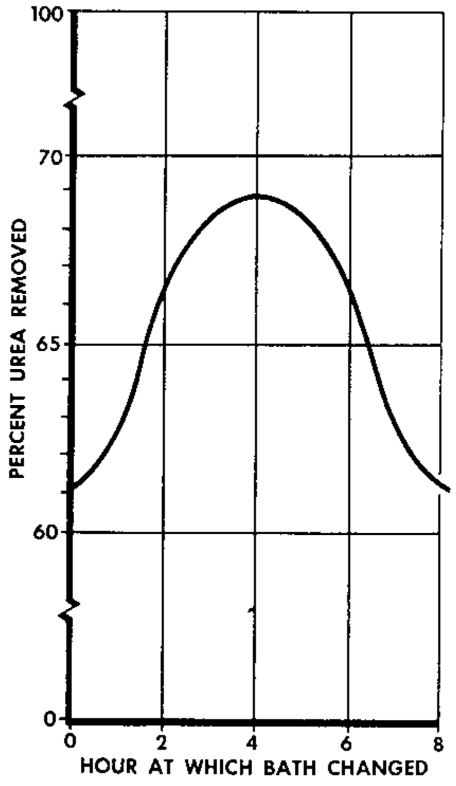

Optimal time to change bath

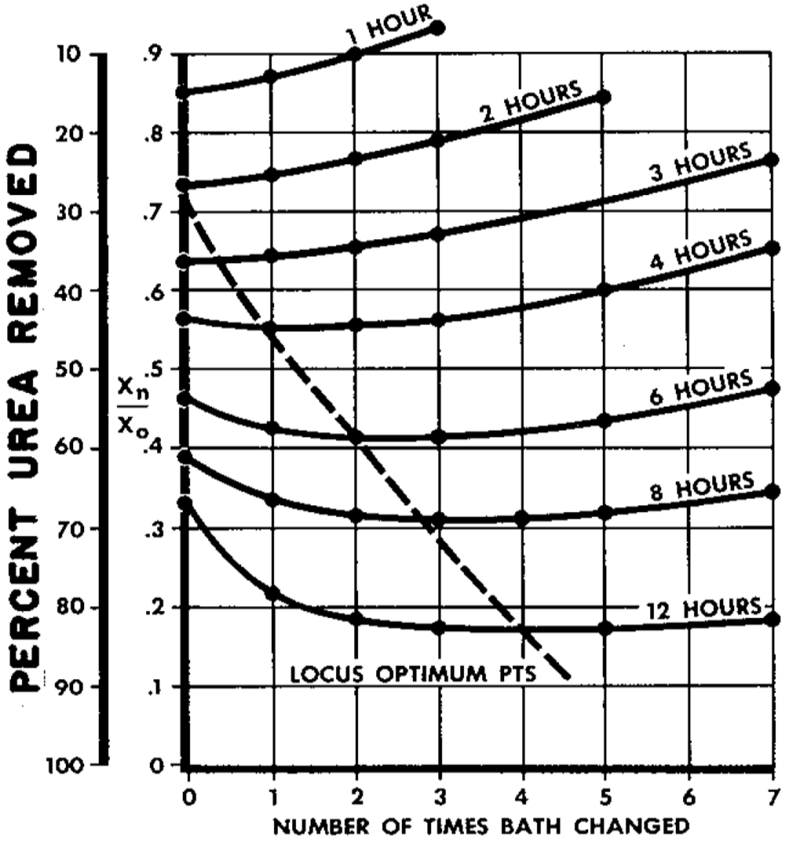

Optimization of Kidney Transplants

Kidney Exchange

Donors

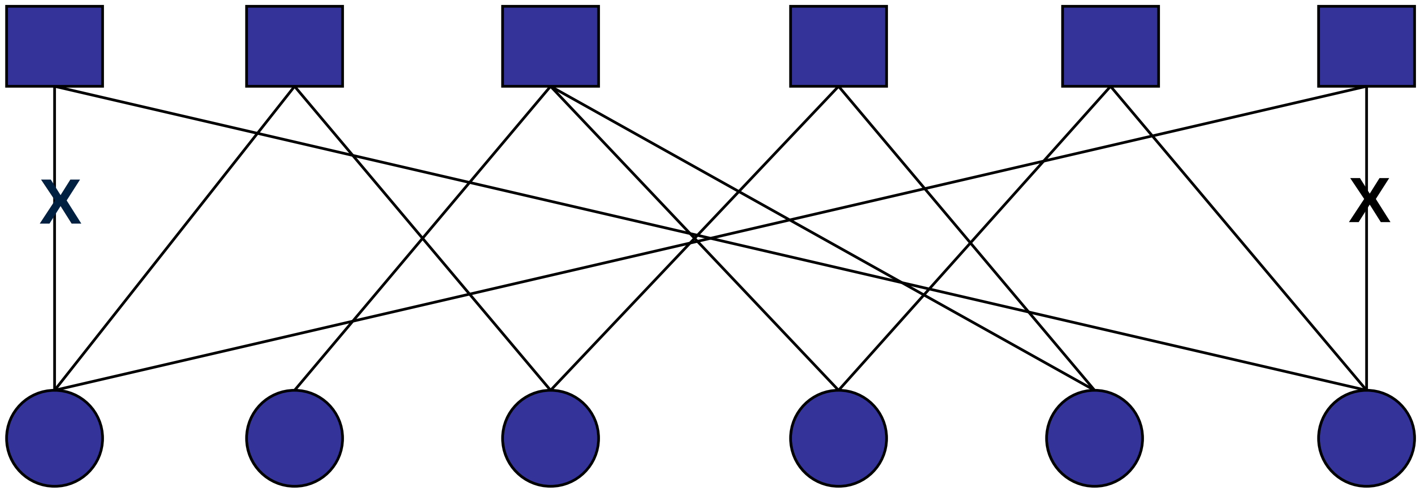

Recipients

Paired Matching

Figure 1. Graph Theory Model of Donor/Recipient Nodes, With Links Indicating Compatible Matches

Segev, D, Gentry, S.E., Warren, D.S, Reeb, Montgomery, RA, 2005, Kidney Paired Donation and Optimizing the Use of Live Donor Organs, *JAMA,* 293(15), 1883-1890.

Criteria

- Number of matches
- Number of priority matches
- Immunologic concordance
- Travel requirements

Constraints

- Compatibility is determined by two primary factors:
  - Blood type
  - Tissue antibodies

- Blood type compatibility

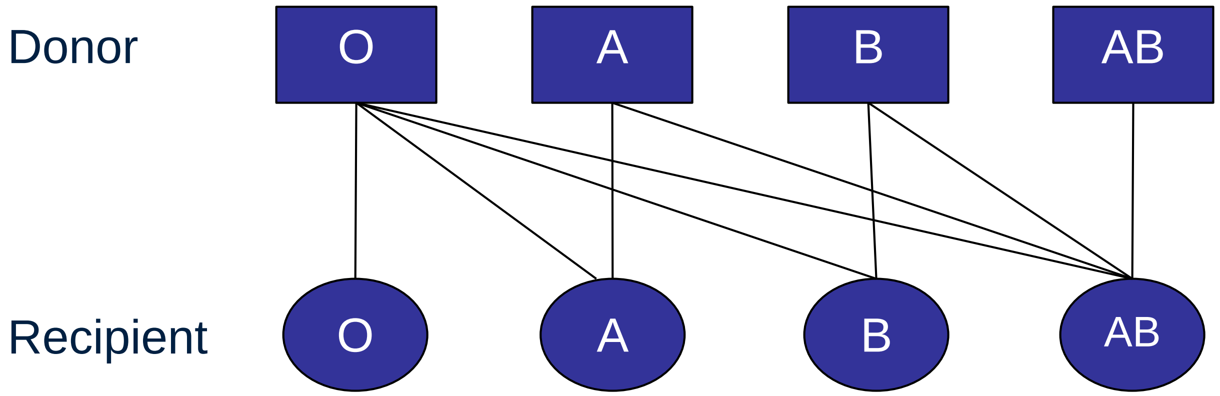

Matching Problems

Given a graph $G(V,E)$ a *matching* is a set of pairwise nonadjacent edges.

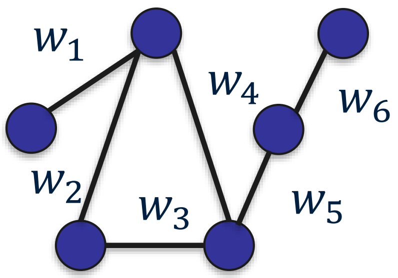

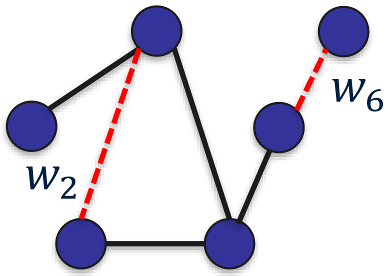

2 matches

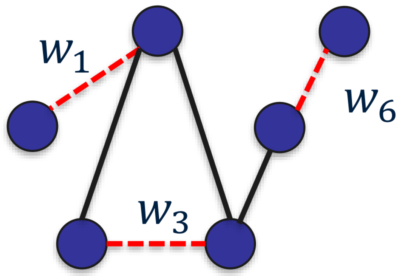

3 matches

A *maximal edge-weight matching* is a set of non-adjacent edges with maximum total weight among all matches.

Maximum Edge Weight Matching

A matching problem for a graph $G(V, E)$ can be expressed as an *integer program*

$$Max \ \ \sum_{e \in E} w_e x_e$$

$$Subject\ to:$$

$$\sum_{e \sim v} x_e \leq 1, \quad for\ all\ v \in V$$

$$x_e \in \{0,1\}, \ for\ all\ e \in E$$

Edmonds, J. 1965. "Paths, trees, and flowers," *Canadian J. Math*. 17: 449–467.

Properties of Matching Graphs

Analysis of factors that influence vertex and edge weights

- In a vertex weighted graph with positive weights any matching with maximum vertex weight has maximum cardinality
- In a matching with maximum edge weight could have **half as many** edges as a maximum cardinality matching
  - The ratio can be bounded by controlling : $\max_i w_i - \min_i w_i$
- Connections to multi-criteria problems:
    - Weighted objectives
    - Bi-level optimization

Gentry, S., Michael, T.S., Segev, D. "Maximum Matching in Graphs for Allocating Kidney Paired Donation," Working Paper

Impact

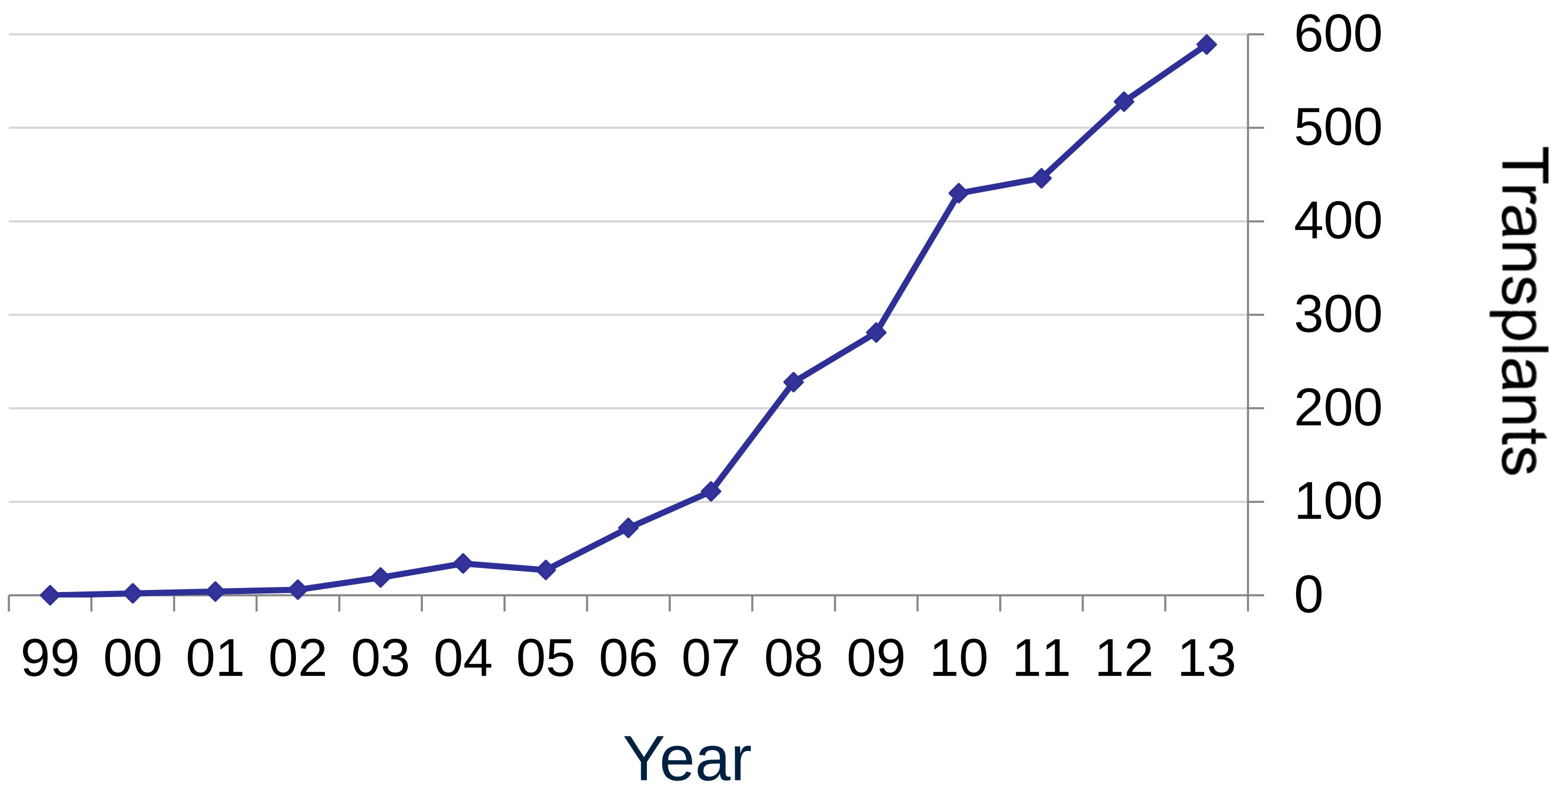

From 1 in 1999, to nearly 600 in 2013, KPD now comprises 10% of living kidney donations**

*Figure courtesy of Sommer Gentry, US Naval Academy; www.optimizedmatch.com

Example 3: Diabetes

- 29 million people have diabetes in the U.S.
  - 9% of the U.S. population
  - 90% have type 2 diabetes
- Health complications include micro and macro-vascular events
- Medication can control major risk factors like blood sugar, cholesterol and blood pressure.

Mason, J.E. et al. 2014. Optimizing the Simultaneous Management of Blood pressure and Cholesterol for Type 2 Diabetes Patients. European Journal of Operational Research; 233(3) 727-738.

Sequential Decision Making

- Choose the best action each time period to maximize long term expected rewards

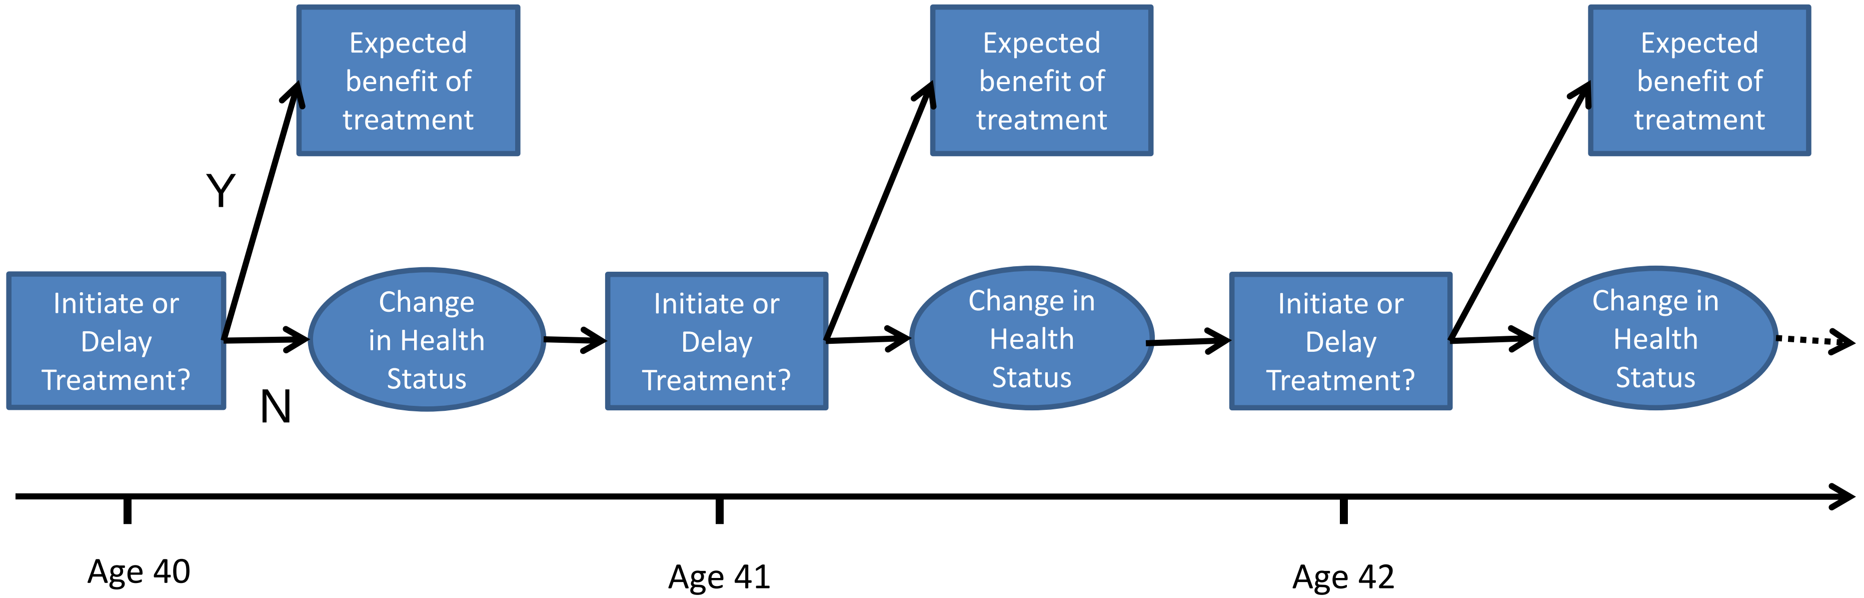

State Transition Diagram

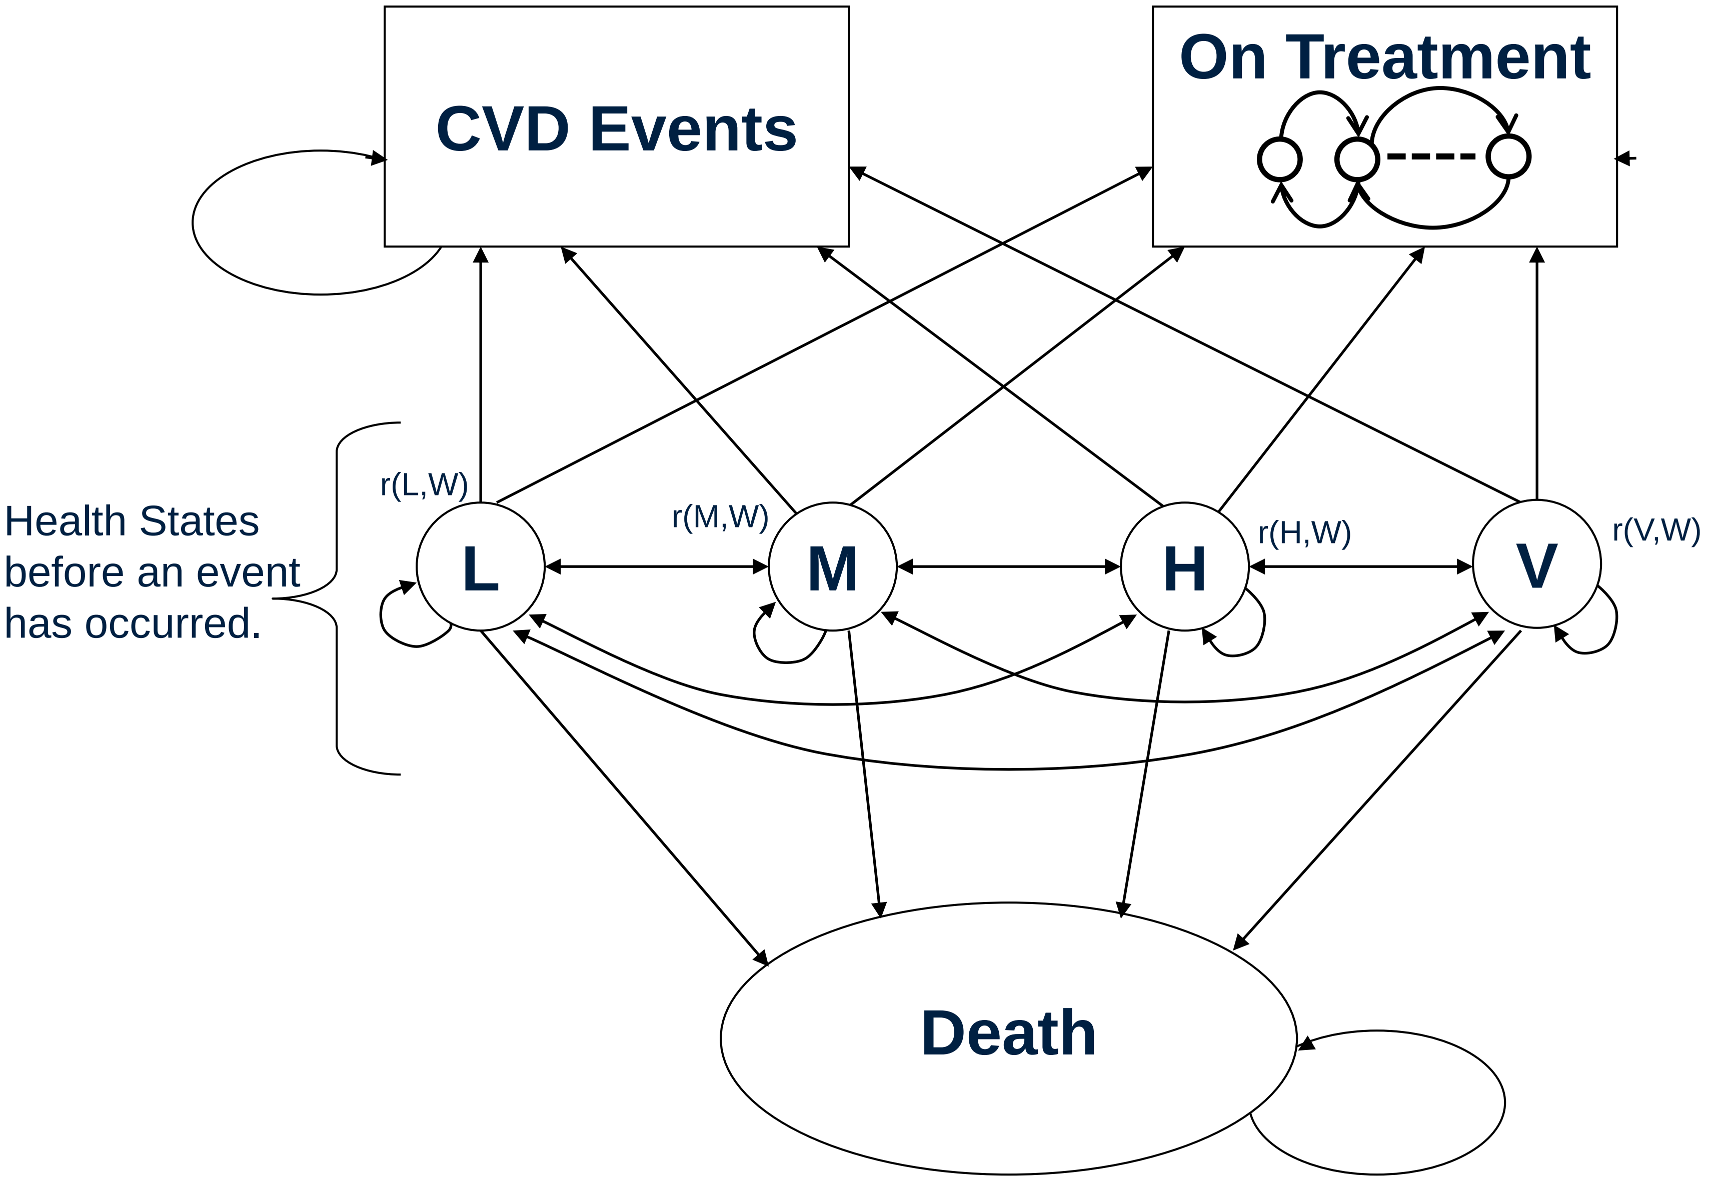

Markov Decision Process

- Health status: $s_t \in S \equiv \{1,2,3, \ldots \ldots L, L+1\}$
- Treatment decision in state s_t: $a(s_t) \in A(s_t)$
- Optimality Equations for all $s_t, t = 1, \ldots, T-1$:

$$\overbrace{v_t(s_t)}^{\text{Optimal Reward to Go in Health State } s_t} = \max_{a_t}\{\overbrace{r(s_t, a_t)}^{\text{Period } t \text{ Reward}} + \lambda \overbrace{\sum_{\forall s_{t+1}} \underbrace{p(s_t' | s_t, a_t)}_{\text{Transition probabilities}} v_{t+1}(s_t')}^{\text{Discounted Expected Future Reward}}\}$$

$$\underbrace{v_T(s_T) = r(s_T)}_{\text{Boundary condition}}$$

Reward Function

Rewards for each state action pair define the objective function for a Markov decision process

$$r(s_t, a_t) = \overbrace{\alpha L(s_t, a_t)}^{\text{Reward for living disease free}} - (1-\alpha)\big(\overbrace{C^S(s_t) + C^{CHD}(s_t)}^{\text{Cost of CVD Events}} + \underbrace{C^M(s_t)}_{\text{Medication Cost}}\big)$$

Policy Evaluation

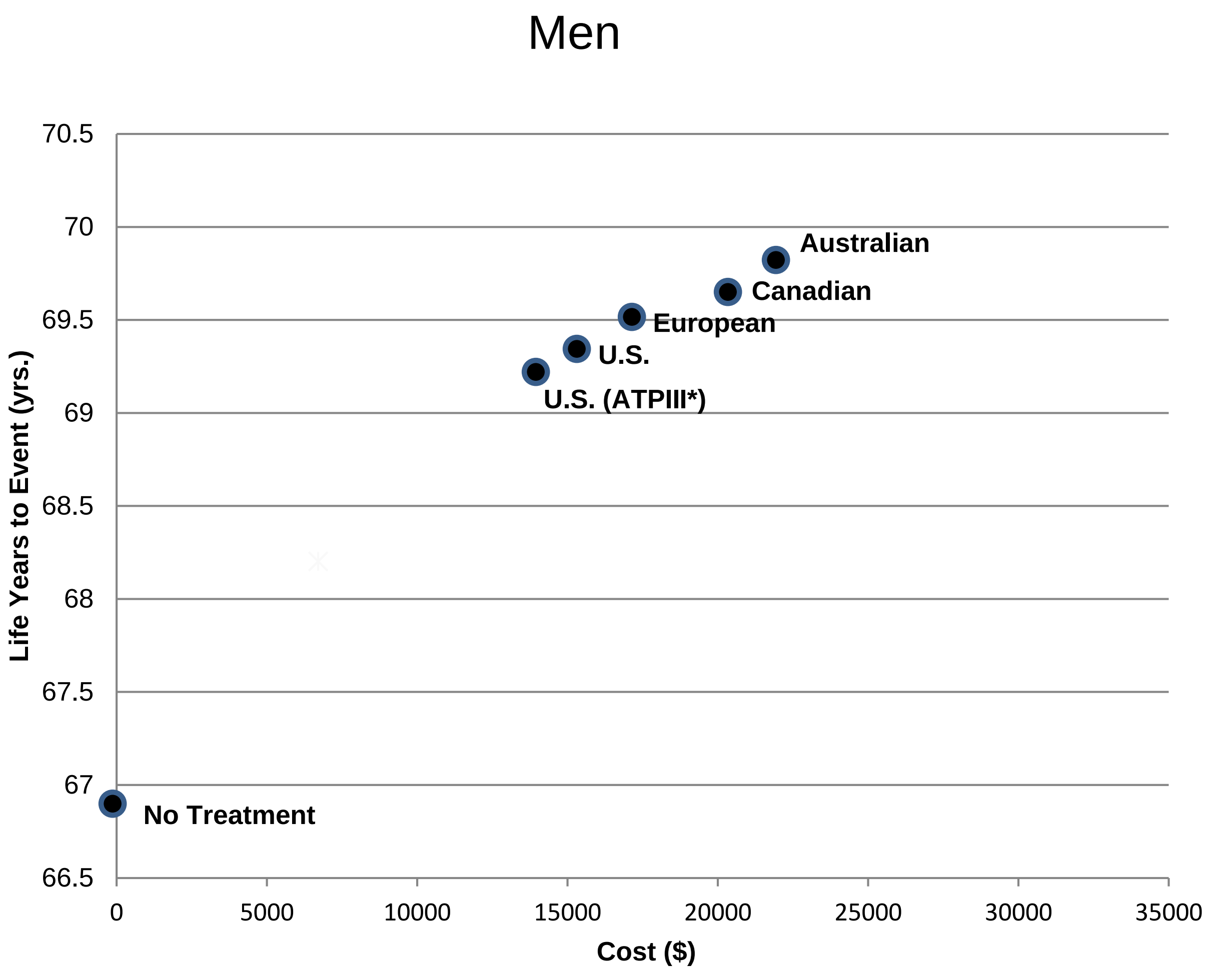

Optimal Policy vs Guidelines

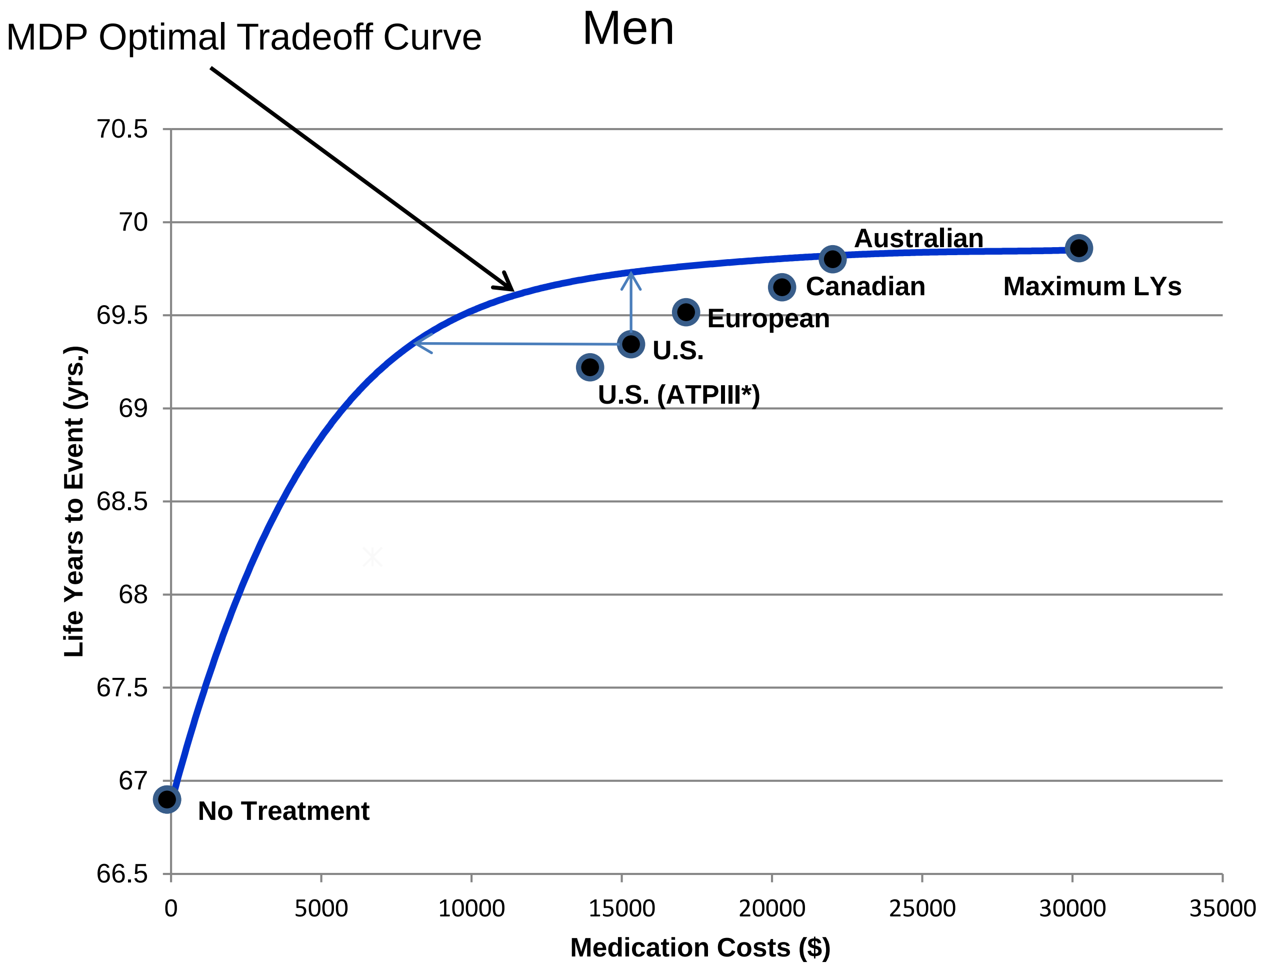

Optimal Policy vs Guidelines

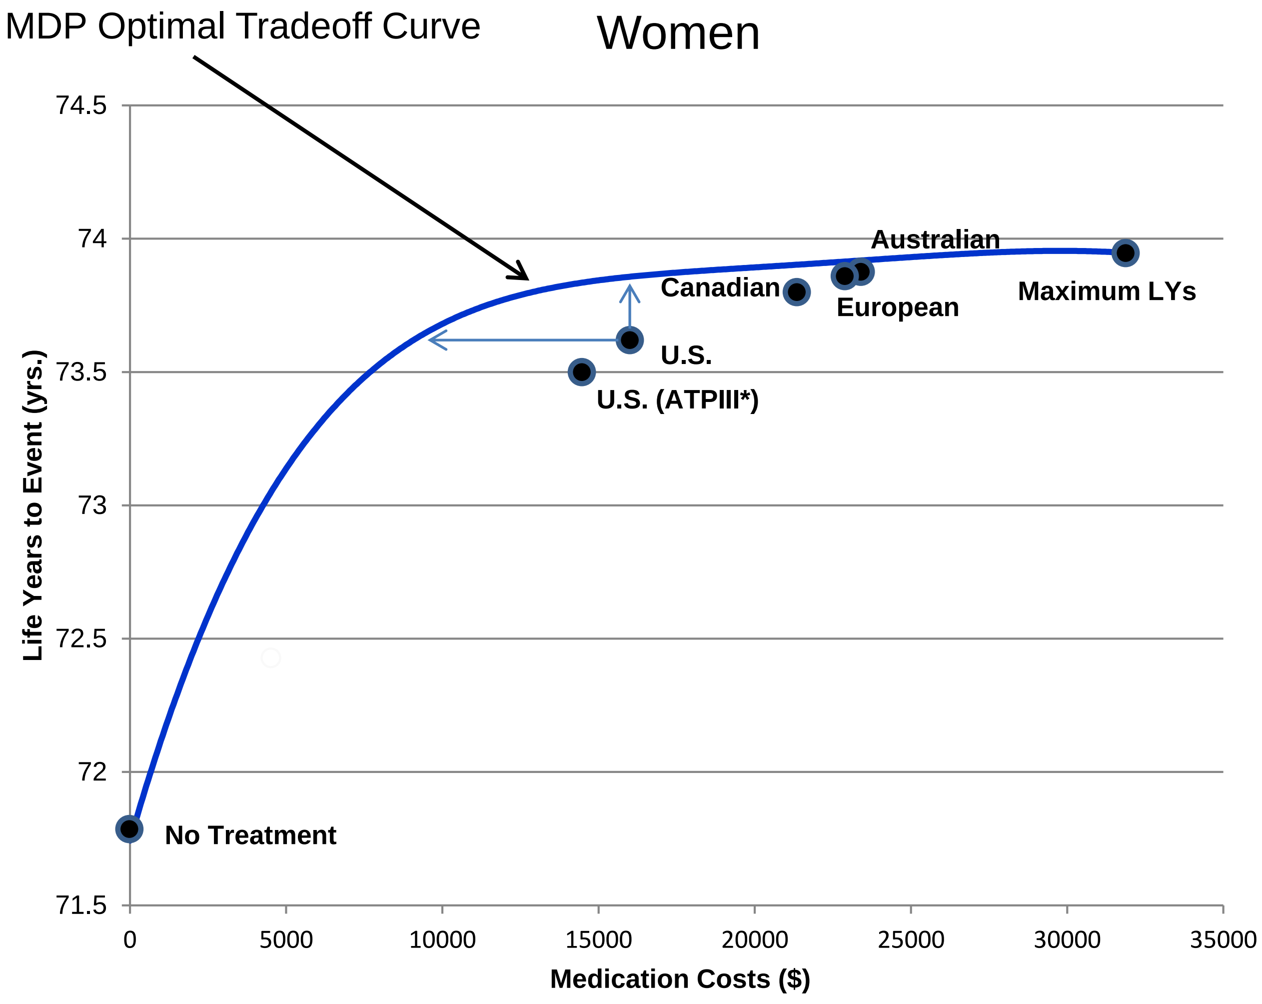

Recent Work: Robust MDPs

- All models are subject to uncertainty in model parameter estimates and model assumptions
    - Transition probabilities are based on statistical estimates from longitudinal data
    - Rewards are based on estimates of mean patient utility, cost, or other performance measures
- Robust MDPs (RMDPs) attempt to account for this uncertainty

Delage, E., Iancu, D. 2015. “Robust Multistage Decision Making.” INFORMS Tutorials in Operations Research

Robust MDPs

- Goal of a standard finite horizon MDP is to find π^* with respect to a fixed Markov chain with TPM, P:

$$\pi^* = \operatorname*{argmax}_{\pi \in \Pi} E^P[\sum_{t=1}^{N-1} r_t(s_t, \pi(s_t)) + r_N(s_N)]$$

- An RMDP can be viewed as a sequential game against an **adversary**:

$$\pi^* = \operatorname*{argmax}_{\pi \in \Pi} \min_{P \in U} \ E^P[\sum_{t=1}^{N-1} r_t(s_t, \pi(s_t)) + r_N(s_N)]$$

Time Varying Model

This problem is “easy” when the **rectangularity assumption** is made:

$$U = \prod_{\forall s_t \in S} U(s_t)$$

Under this assumption the optimality equations are:

$$v_t(s_t) = \max_{a_t \in A} \left\{ r_t(s_t, a_t) + \max_{p(s_t) \in U(s_t)} \lambda \sum_{s_{t+1} \in S} p(s_{t+1} | s_t, a_t)\, v_{t+1}(s_{t+1}) \right\}$$

Where $p(s_t)$ is the row of the TPM corresponding to state s_t and $U(s_t)$ is the row’s uncertainty set.

RMDP Case Study: Type 2 Diabetes

Many medications that vary in efficacy, side effects and cost.

Oral Medications:
- Metformin
- Sulfonylurea
- DPP-4 Inhibitors

Injectable Medications:
- Insulin
- GLP-1 Agonists

Zhang, Y., Steimle, L.N., Denton, B.T., 2017. “Robust Markov Decision Processes for Medical Treatment Decisions,” Working Paper, available at *Optimization Online*.

Treatment Goals

- HbA1C is an important biomarker for blood sugar control
- But disagreement exists about the optimal goals of treatment and which medications to use

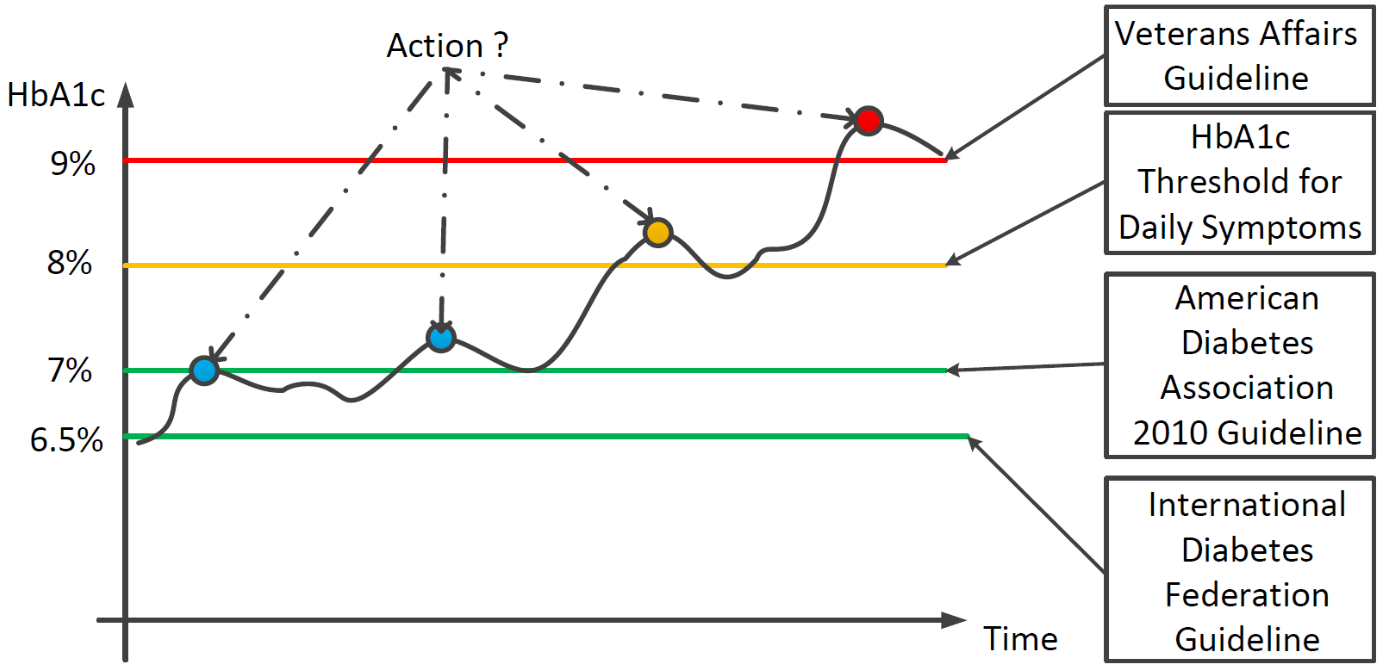

Markov Chain for Type 2 Diabetes

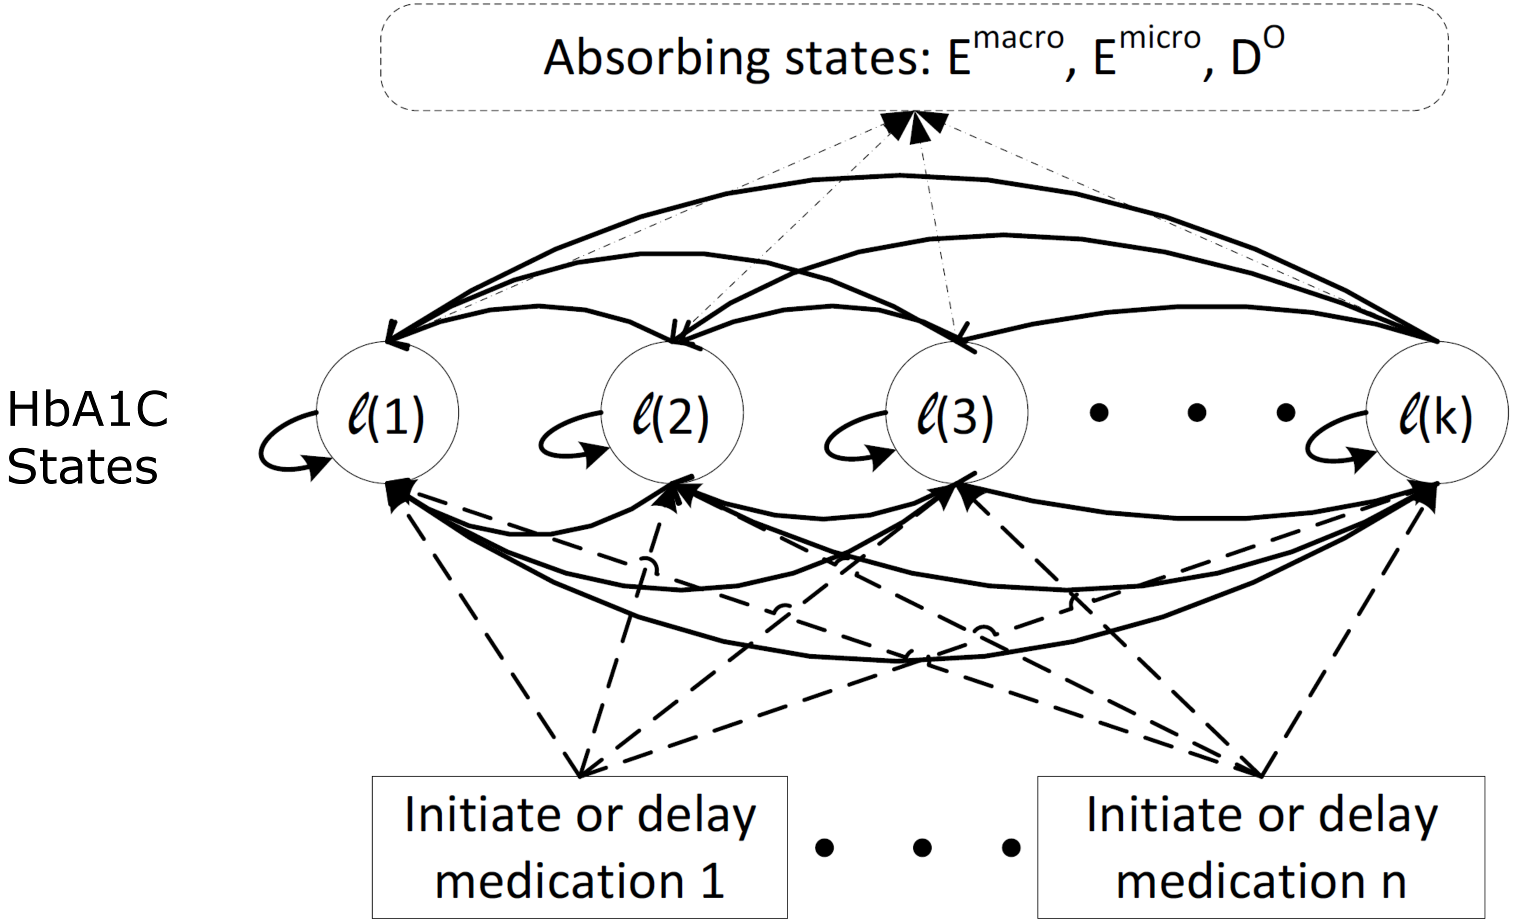

Uncertainty Set with Budget

$$U(s_t) = \begin{cases} p(s_{t+1}|s_t) = \hat{p}(s_{t+1}|s_t) - \delta^L z^L(s_{t+1}) + \delta^U z^U(s_{t+1}), & \forall s_{t+1} \\ \sum_{s_{t+1} \in S} p(s_{t+1}|s_t) = 1 & \\ \boxed{\sum_{s_{t+1}} (z^L(s_{t+1}) + z^U(s_{t+1})) \leq \Gamma(s_{t+1})} & \\ z^L(s_{t+1}) \cdot z^U(s_{t+1}) = 0, & \forall s_{t+1} \\ 0 \leq p(s_{t+1}|s_t) \leq 1, & \forall s_{t+1} \end{cases}$$

Properties:

- Can be reformulated as a linear program
- For $\Gamma = |S|$ can be solved in $O(|S|)$

Uncertainty Set

A combination of laboratory data and pharmacy claims data was to estimate transition probabilities between deciles

$$p(s'|s\ ,a) = \frac{n(s,s',a)}{\sum_{s'} n(s,s',a)}\,, \forall s', s, a$$

$1 - \alpha$ confidence intervals for row s of the TPM:

$$[\hat{p}(s'|s,a) - S(\hat{p}(s'|s,a)L, \quad \hat{p}(s'|s,a) + S(\hat{p}(s'|s,a)L]$$

where

$$S(\hat{p}(s'|s,a)L = \left[\chi^2_{|s|-1,\alpha/2|S|)} \frac{\hat{p}(s'|s,a)\big(1-\hat{p}(s'|s,a)\big)}{N(s)}\right]^{\frac{1}{2}}$$

Gold, Ruth Z. "Tests auxiliary to χ 2 tests in a Markov chain." *The Annals of Mathematical Statistics* 34, no. 1 (1963): 56-74.

Results

Quality adjusted life years to first health complications for women with type 2 diabetes

Simulated

Worst-case

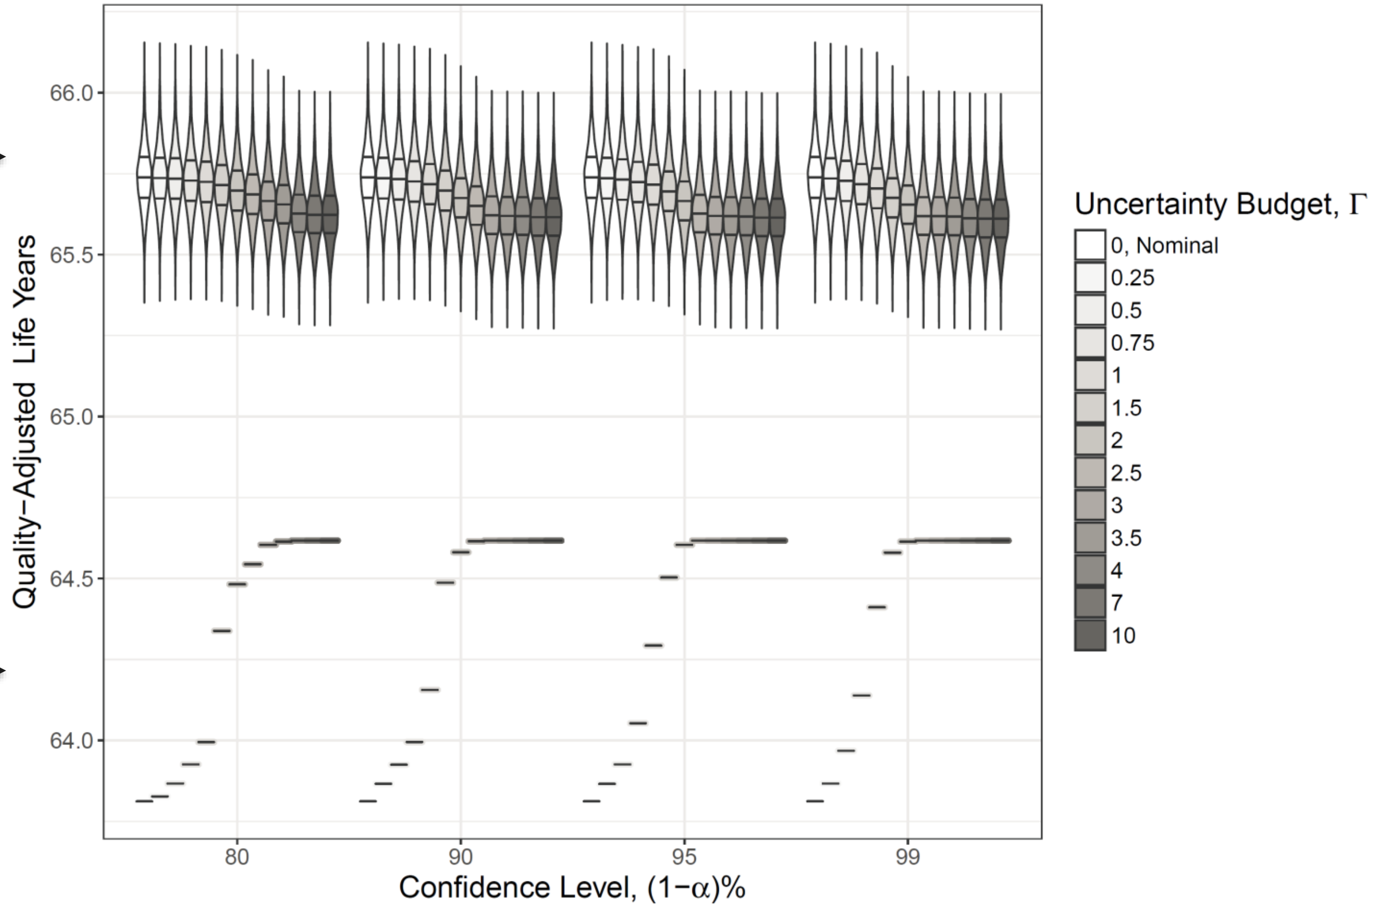

Results

Mean QALYs versus variance in QALYs to first event for women with type 2 diabetes

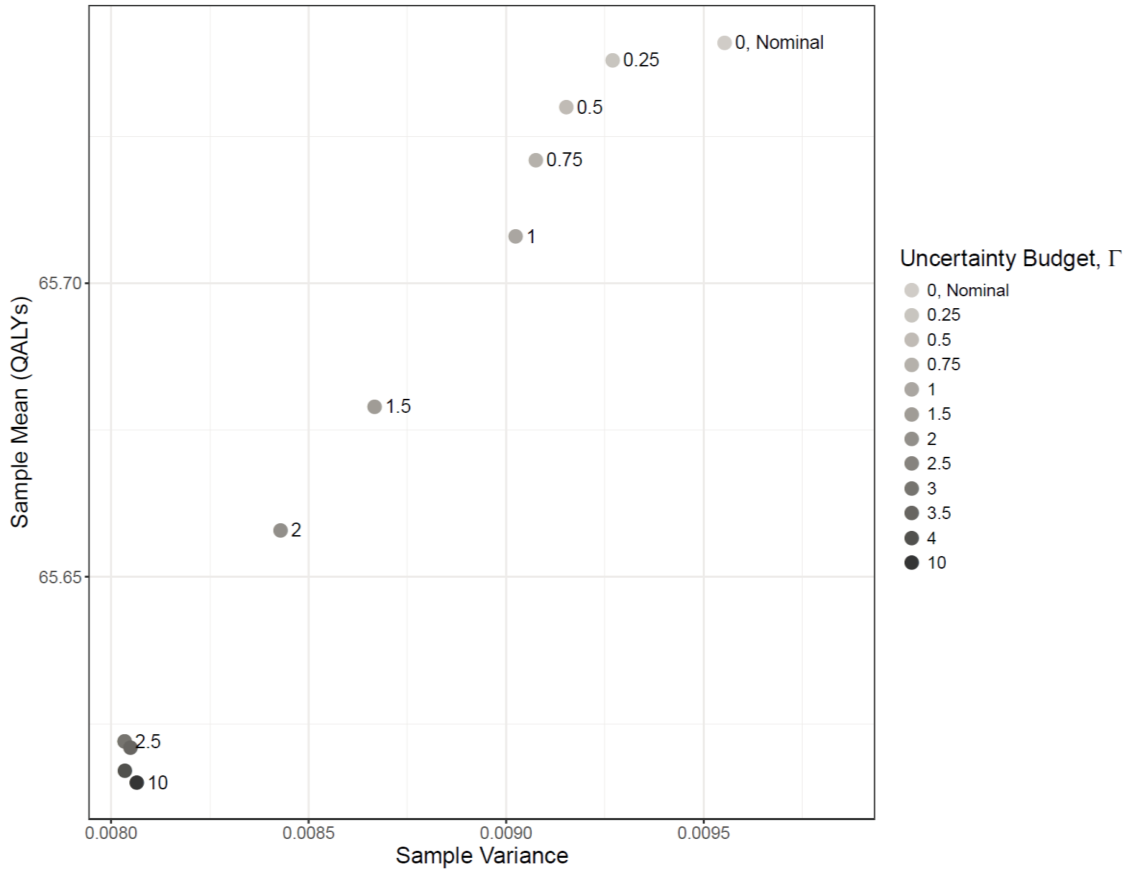

Other Work

- Liver Transplants: Alagoz, Maillart, Schaefer, Roberts, *Management Science*, 2004
- Breast Cancer: Maillart, Ivy, Ransom, Dielhl, *Operations Research*, 2008
- HIV: Shechter, Schaefer, Roberts, *Operations Research*, 2008
- Prostate Cancer: Zhang, Denton, Balasubramanian, Shah, *M&SOM* 2012
- Adherence to Screening: Ayer, Alagoz, Stout, Burnside, Management Science, 2015
- Colorectal Cancer: Erenay, Alagoz, Said, *M&SOM* , 2014

Markov Decision Processes for Screening and Treatment of Chronic Diseases

Markov Decision Processes in Practice pp 189-222

Part of the International Series in Operations Research & Management Science book series (ISOR, volume 248)

- Lauren N. Steimle (1) Email author (steimle@umich.edu)
- Brian T. Denton (1)

1. Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, USA

Optimization in Medicine:
The Future

PubMed Search Results

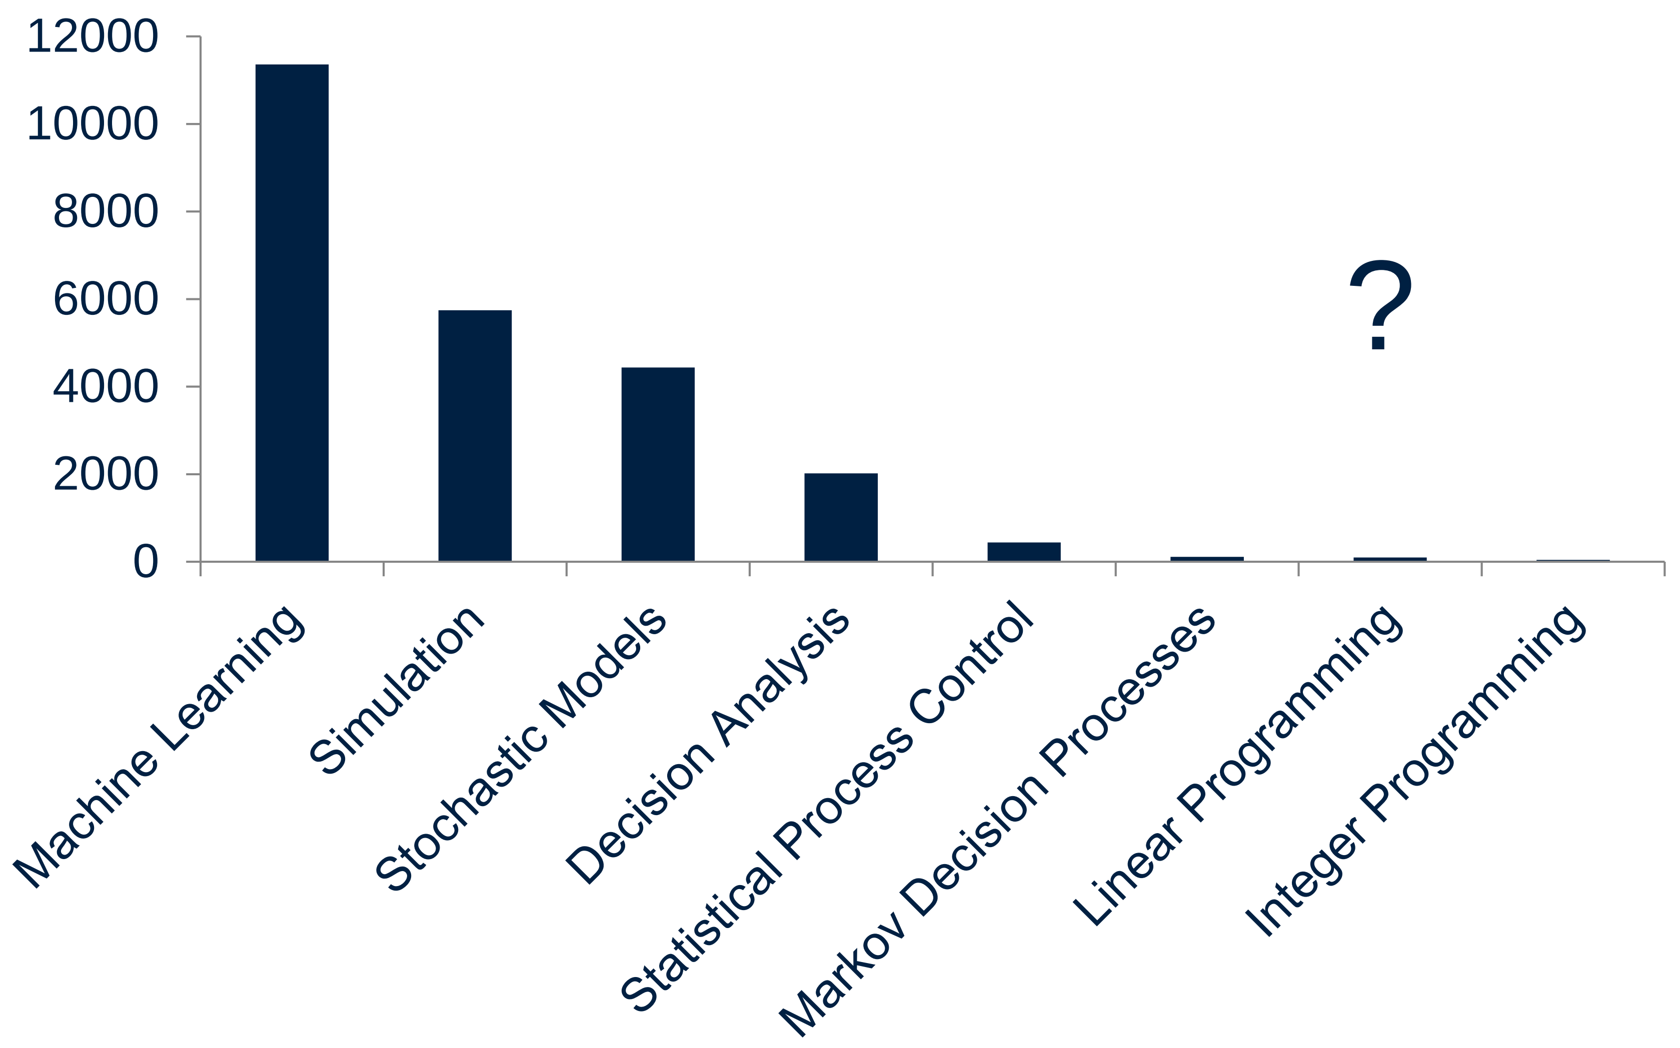

Linking Decisions Across Time

Research Questions:

- When and how to screen for diseases?
- When to use diagnostic tests?
- When to treat?
- When to stop?

Methods for Sequential Decision Making:

- Decision Analysis
- Markov decision processes
- Partially observable Markov decision processes
- Multi-stage stochastic programming
- Reinforcement learning

Example

UVA's Continuous Closed-Loop Artificial Pancreas Powered by Android Smartphone

- Difficult real time optimal control problem
- Must maintain glucose levels within a defined range
- Current glucose state difficult to predict

Cobelli, C, Renard,E., Kovatchev, B. 2011. Artificial Pancreas: Past, Present, Future, *Diabetes*, 60, 2682 - 2682

Resource Constrained Decision Making

Research Questions:

- Coordination across medical “silos”
- Prioritizing treatment in resource constrained settings:
  - High value health care
  - Constrained burden on patients

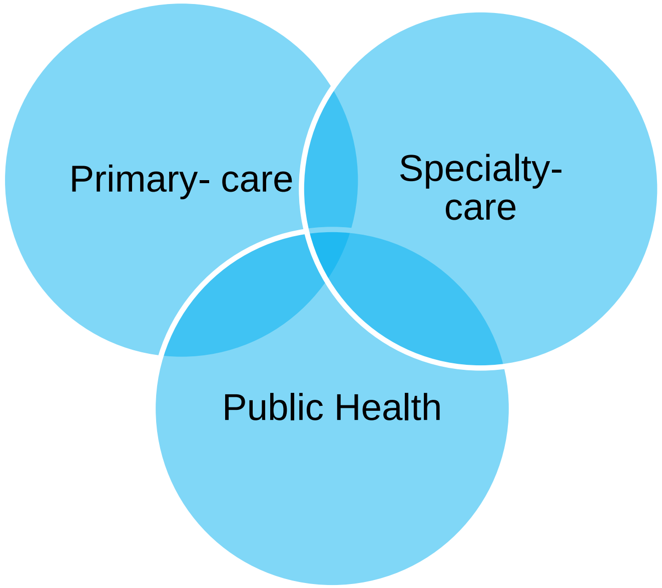

Methods:

- Machine Learning
- Mathematical programming
- Sequential decision making
- Intuitive approximations

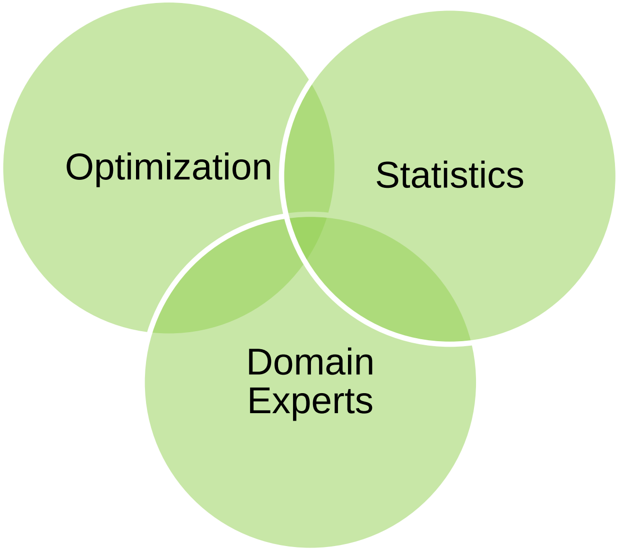

Example

Optimizing the coordination of imaging decision for cancer staging

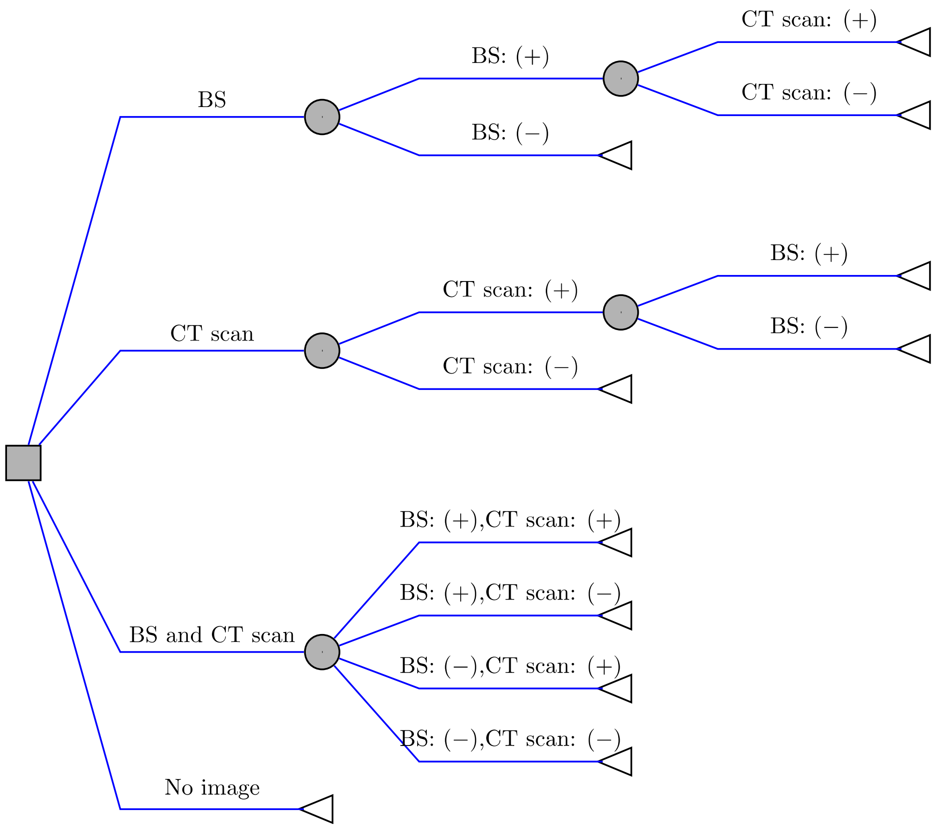

Statistical error

$$\min \sum_{\forall j} \sum_{\forall p} n_{kp} z_{jp}$$

Subject to:

$$\sum_{\forall j} \sum_{\forall p} m_{jp} z_{jp} \leq \alpha$$

$$\sum_{\forall p} z_{jp} = 1, \qquad \forall j$$

$$z_{jp} \in \{0,1\}, \quad \forall j, p$$

Merdan, S., Barnett, C., Denton, B.T., Montie, J.E., Miller, D.C., 2017. "Data Analytics for Optimal Detection of Metastatic Prostate Cancer" Working Paper

{Optimization} ∩ {Statistics}

$$max \ \{ f(x) \mid x \in C \}$$

Data

Missing Data

Bias

Causal Inference

Uncertainty and Ambiguity

Takeaways

- Optimization can improve medical decision making and vice versa but…

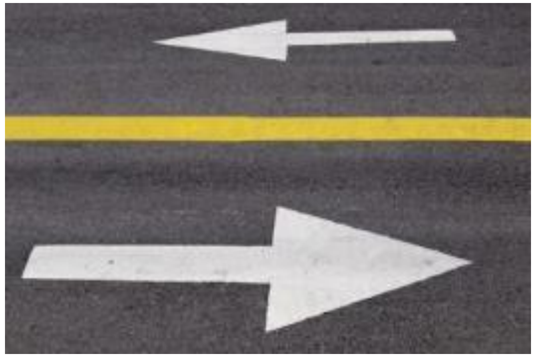

- It is underutilized and there are many challenges and unexplored opportunities to address this problem

Acknowledgements

Christine Barnett, University of Michigan
Marina Epelman, University of Michigan
Sommer Gentry, US Navel Academy
Jennifer Mason, University of Virginia
Selin Merdan, University of Michigan
Lauren Steimle, University of Michigan
Edwin Romeijn, GA Tech
Nilay Shah, Mayo Clinic

Some of this work was funded in part by grants from the *CMMI Division* at the *National Science Foundation*

COLLEGE OF ENGINEERING
INDUSTRIAL & OPERATIONS ENGINEERING
UNIVERSITY OF MICHIGAN

Brian Denton
Industrial and Operations Engineering
University of Michigan

btdenton@umich.edu

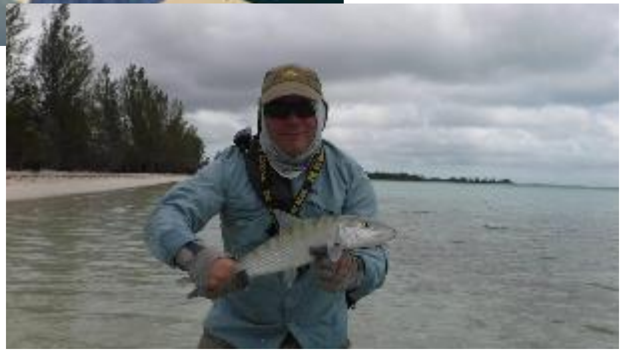

These slides (and pictures!) are on my website:

https://btdenton.engin.umich.edu/presentations/