Constructing stabilized dynamic treatment regimens using electronic health record data

Yingqi Zhao

Fred Hutchinson Cancer Research Center

Oct 25, 2016
2nd Seattle Symposium on Health Care Data Analysis

Outline

Diabetes

- Current diabetes guidelines: tight control of glycosylated hemoglobin (A1c) (< 7 %)
  - Healthy patients.
  - Based on trials of younger patients without severe diabetes complications or other comorbidities.
  - Relatively low risk of tight control; significant benefits in reducing incidence of vascular events
  - Tight control of BP ($<$130/80 mm Hg) and LDL cholesterol ($<$100 mg/dl) for patients with diabetes

Diabetes

- Inappropriate for complex diabetes patients, i.e., older patients (age $>$ 65 years) and/or those with comorbid conditions.
  - Evidence for these guidelines was mainly obtained from the results of randomized clinical trials (RCTs)
  - Complex patients usually meet the exclusion criteria of clinical trials
  - Increased risk of drug-related morbidity, e.g., hypoglycemia, hypotension

Diabetes

- Guidelines recognize that less stringent treatment goals may be appropriate for complex diabetes patients, and recommend individualization in treatments based on clinical experiences.
- How can we strengthen the current guidelines for complex patients?
- Opportunities: large electronic health records systems

A1c control observational study (PI: Smith, Maureen)

- Linked claims and EHR data for Medicare beneficiaries in the University of Wisconsin Medical Foundation (UWMF) system.
  - Met a validated algorithm for identifying patients with diabetes via claims (each claim contains information associated with the services or procedures performed, e.g., ICD-9-CM diagnosis codes);
  - Medically homed at the participating large, Midwestern, multi-specialty provider group
  - UWMF EHR systems: detailed clinical results including laboratory values and vital signs
- 8,304 diabetes patients active during 2003-2011, recorded each 90-day quarter in which they were alive at the start of the quarter

Data

- A1c values recorded in the EHR.
- Outcome: adverse outcomes occurring during these quarters (e.g., emergency department use or hospitalizations, death), documented from claims information.
- Covariates: sociodemographics, and indicators for comorbidities; time varying patient complexity: the presence of chronic kidney disease (CKD) or congestive heart failure (CHF)

Data

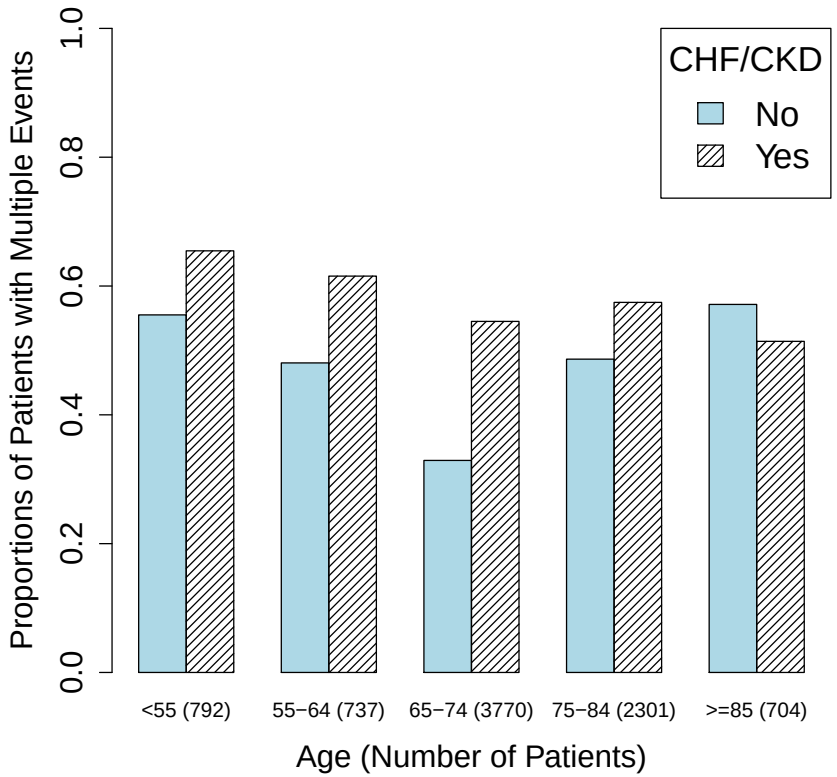

Figure 1: Risk of multiple events.

Data

How to target tight A1c control for these patients?

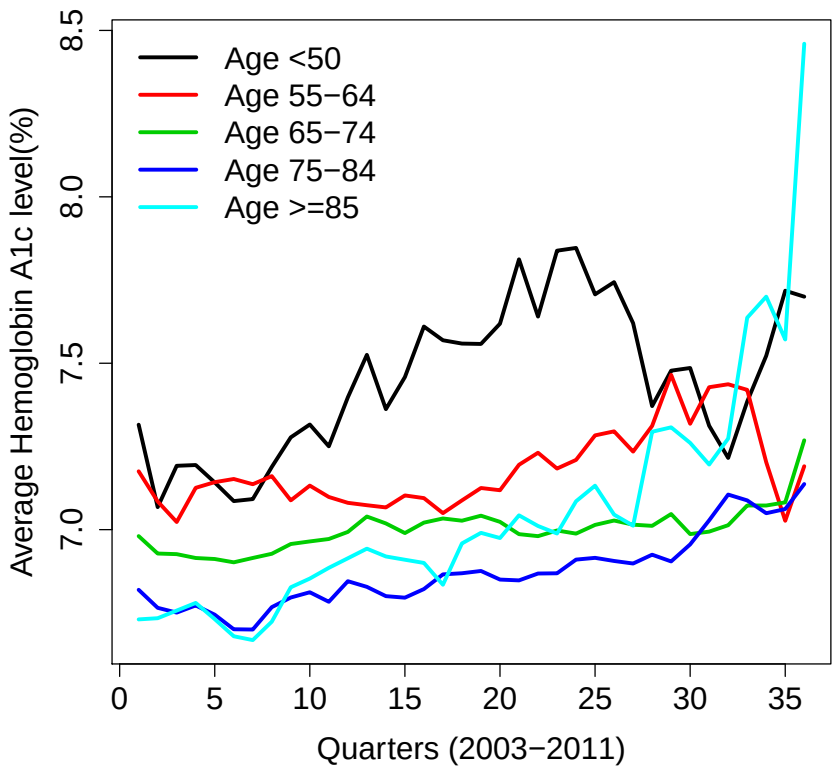

Figure 2: A1c level over time.

Dynamic Treatment Regime

- At any decision point
  - Input: available historical information on the patient to that point.
  - Output: next treatment.

- Dynamic treatment regimes (DTRs) are sequential decision rules for individual patients that can adapt over time to an evolving illness.
  - One decision rule for each time point.
  - Each rule: recommends the treatment at that point as a function of accrued historical information.
  - An algorithm for treating any patient.
  - Aim to optimize some cumulative clinical outcome.

DTR Goals

Learn adaptive treatment strategies: tailor (sequences of) treatments based on patient characteristics.

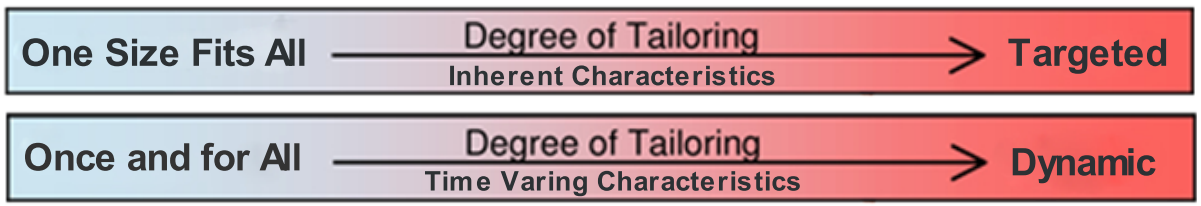

Maximize the benefit of dynamic treatment regimes:

- Well chosen tailoring variables.
- Well devised decision rules.

Dynamic Treatment Regimes (DTRs)

Observe data on n individuals, T stages for each individual,

$$X_1, A_1, R_1, X_2, A_2, \ldots, X_T, A_T, R_T, X_{T+1}$$

- X_1: Initial information.
- X_t: Intermediate information between stages $t-1$ and $t, t \geq 2$.
- A_t: Observed treatment received at stage t, e.g., tight control A1c or not, $A_t \in \{0, 1\}$.
- R_t: Observed outcome following stage t, e.g., ED visits, hospitalizations or deaths following stage t.
- H_t: History available at stage t, $H_t = \{X_1, A_1, R_1, \ldots, A_{t-1}, R_{t-1}, X_t\}$.

A DTR is a sequence of decision rules:

$$d = (d_1(H_1), \ldots, d_T(H_T)), d_t(H_t) \in \{0, 1\}.$$

Multi-stage Data

H_1

Multi-stage Data

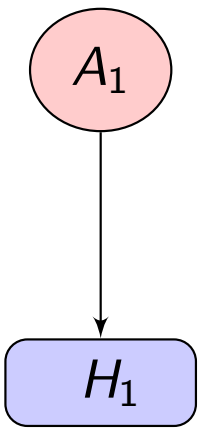

Multi-stage Data

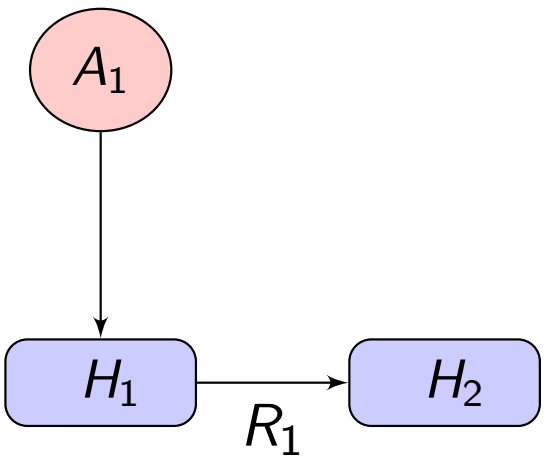

Multi-stage Data

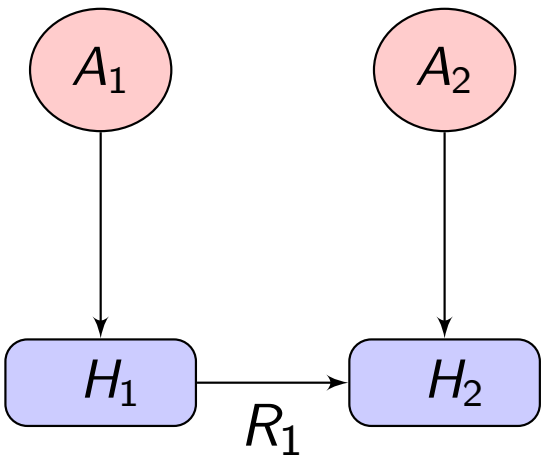

Multi-stage Data

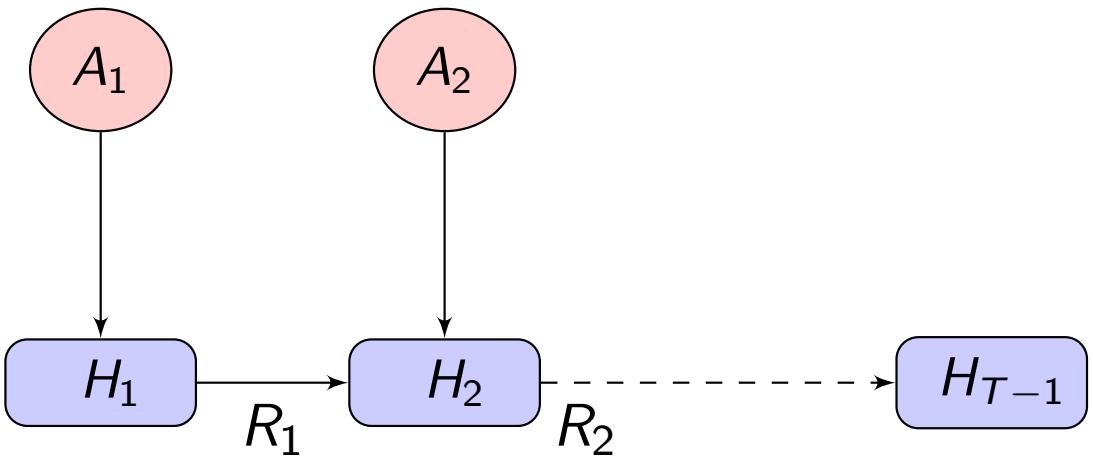

Multi-stage Data

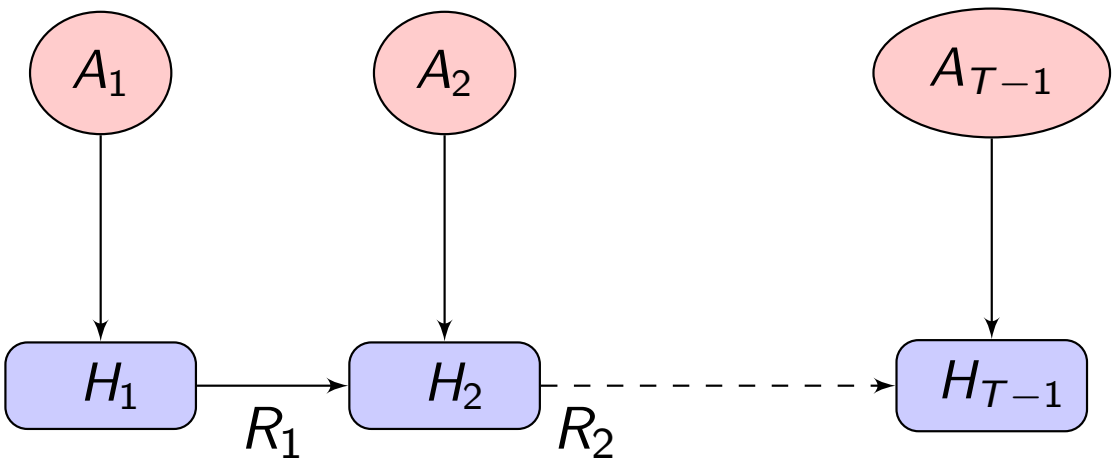

Multi-stage Data

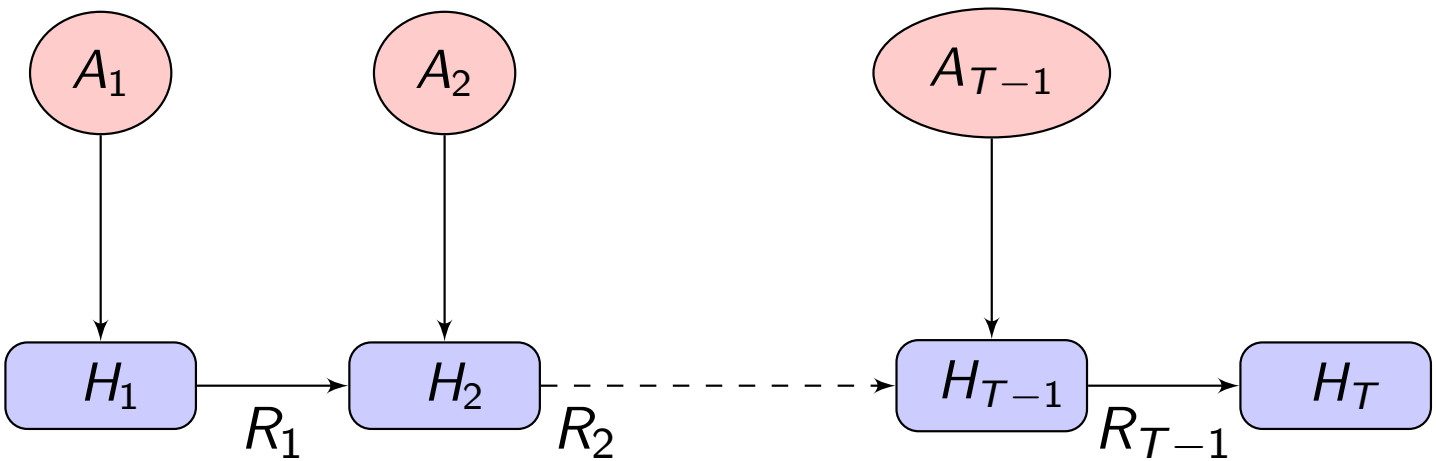

Multi-stage Data

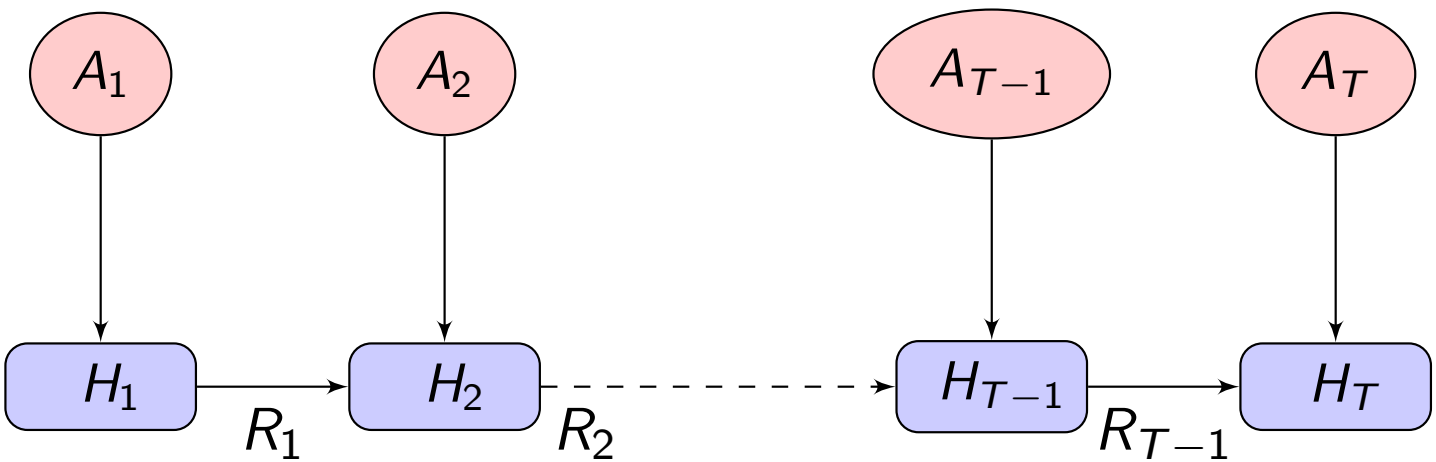

Multi-stage Data

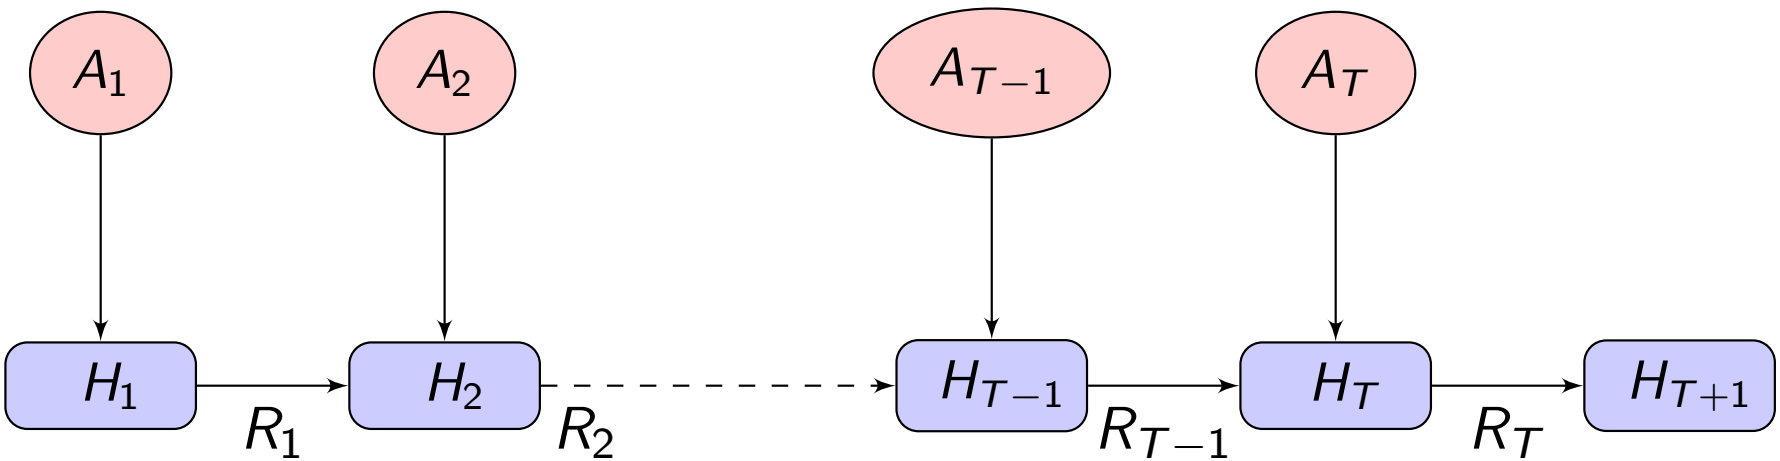

Dynamic Treatment Regimes

An example of desired regime:

Target tight A1c control on patients with A1c level $> 7.5\%$ who does not have CKD/CHF

Mathematically, the formal rule is

$$d_t(H_t) = I(A1c_t > 7.5\% \ \& \ CKD_t = 0 \ \& \ CHF_t = 0).$$

The treatment decision rule at each decision point is the same function of the time-varying covariate(s), also called shared decision (Chakraborty et al (2016))

Value Function and Optimal DTR for Multiple Stages

Goal

Maximize (minimize) the expected sum of outcomes if the DTR is followed by all patients in the population.

- The value function: $\mathcal{V}(d) = E^d(R_1 + \ldots + R_T)$.
- Optimal DTR: $d^* = \text{argmax}_d \, \mathcal{V}(d)$.
- Two main challenges in developing optimal DTRs:
  - Taking individual information into account in decision making.
  - Incorporating long-term benefits and risks of treatment due to delayed effects.

Dynamic Programming

- Estimate d^* if one knows the complete probability distribution of data generation.

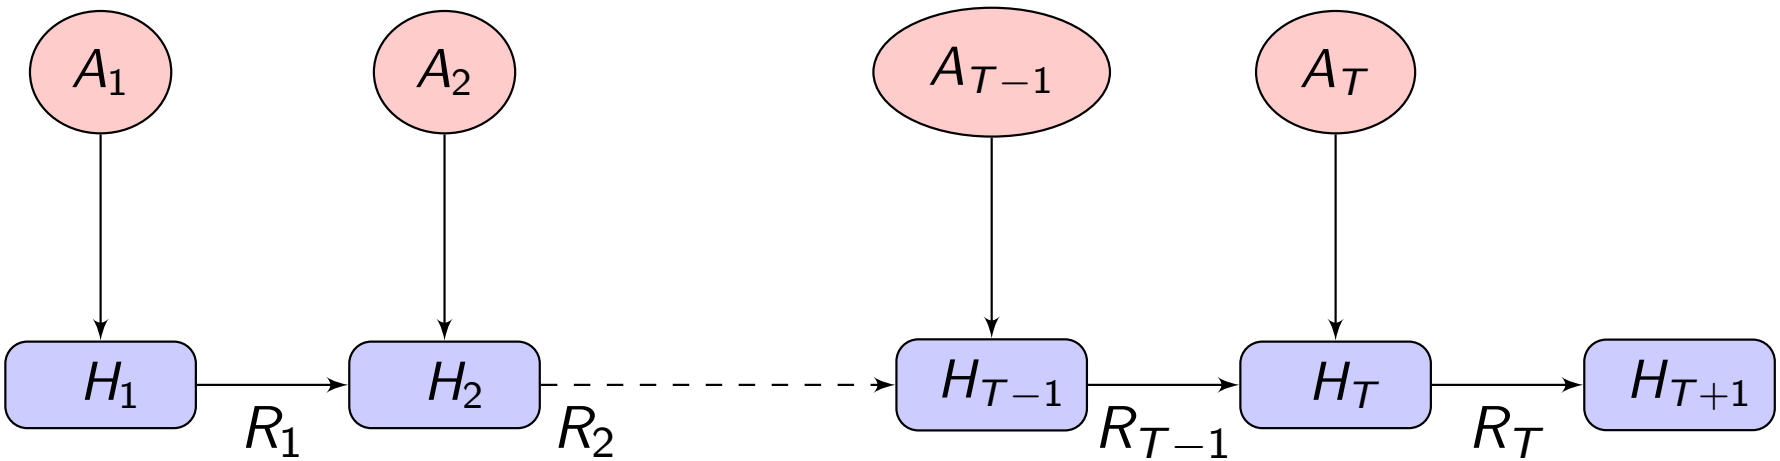

Dynamic Programming

- Estimate d^* if one knows the complete probability distribution of data generation.

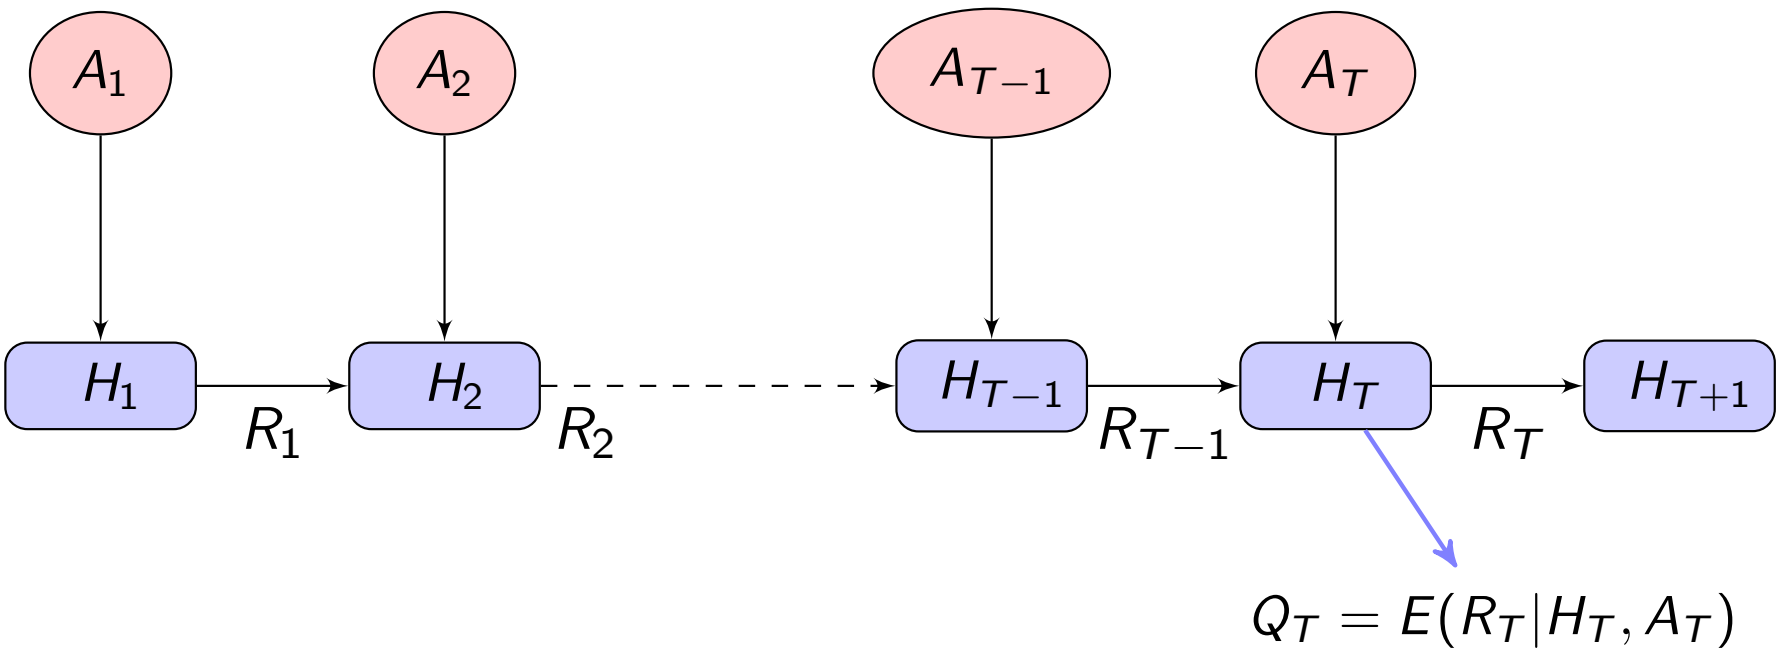

Dynamic Programming

- Estimate d^* if one knows the complete probability distribution of data generation.

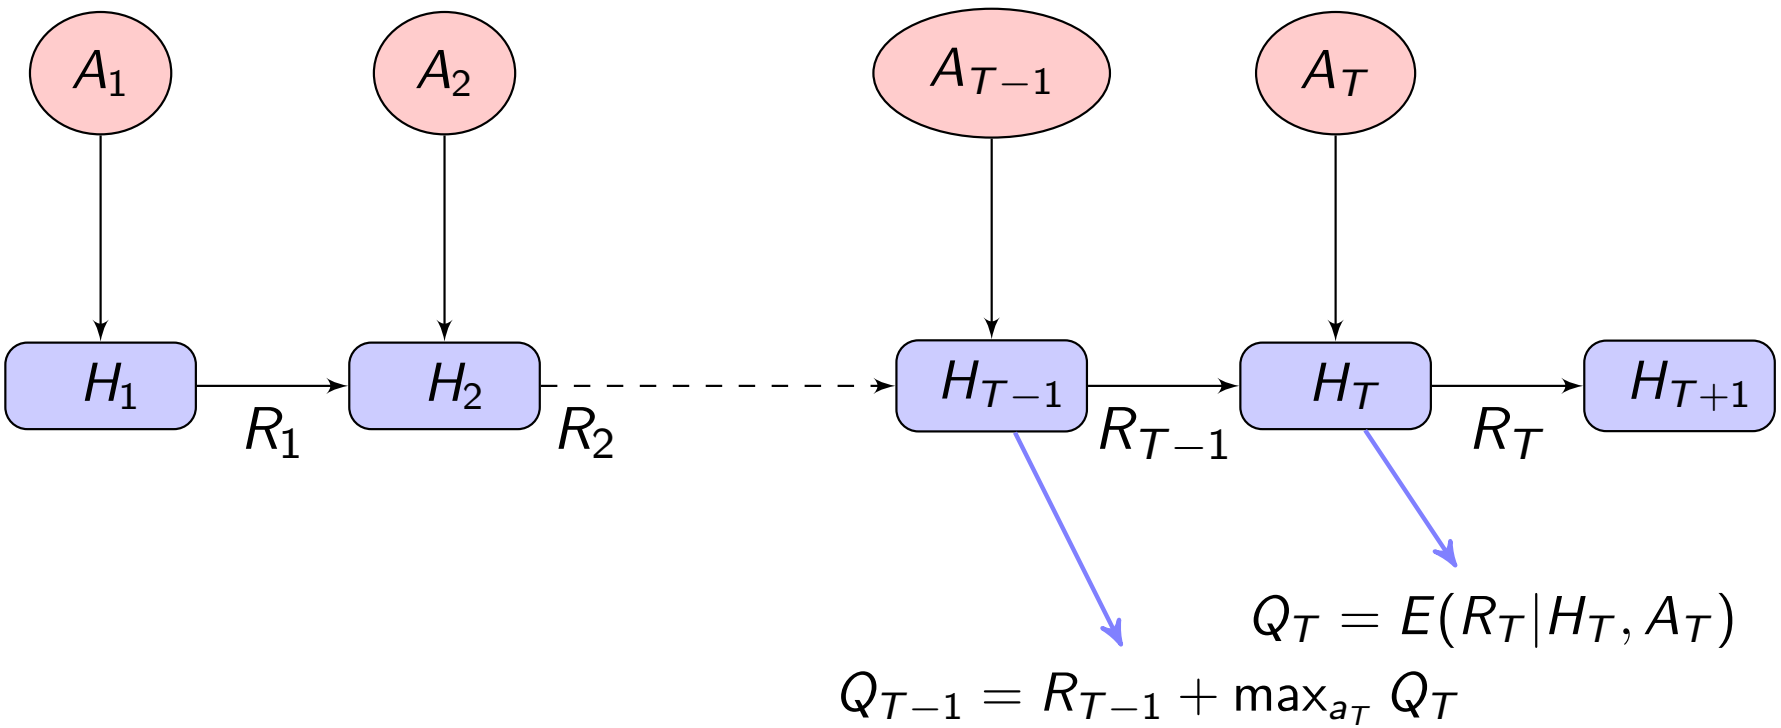

Dynamic Programming

- Estimate d^* if one knows the complete probability distribution of data generation.

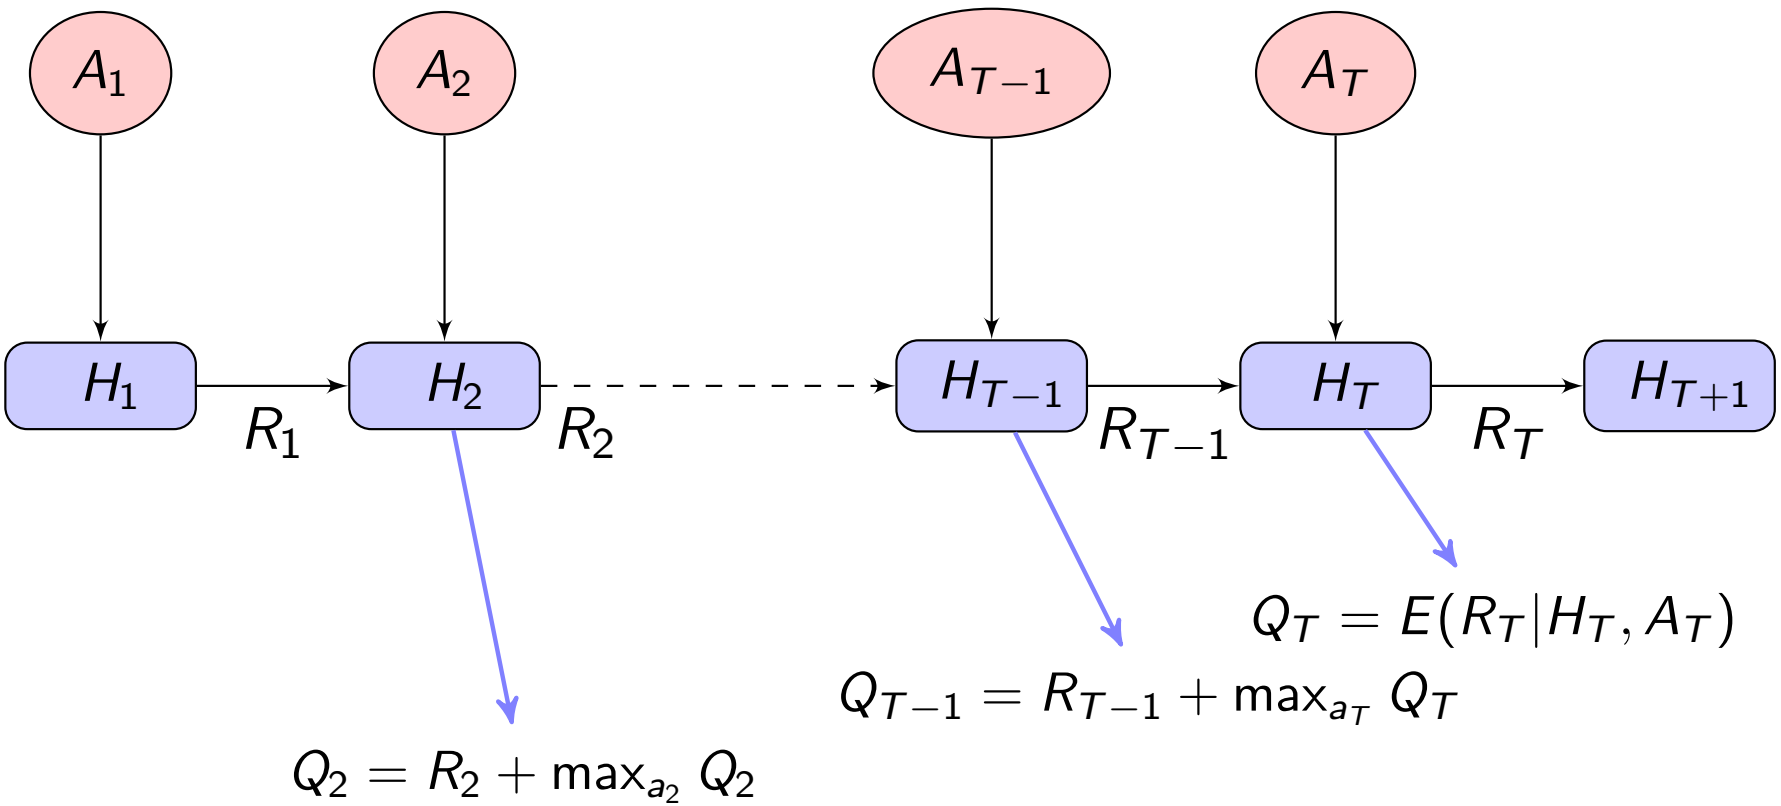

Dynamic Programming

- Estimate d^* if one knows the complete probability distribution of data generation.

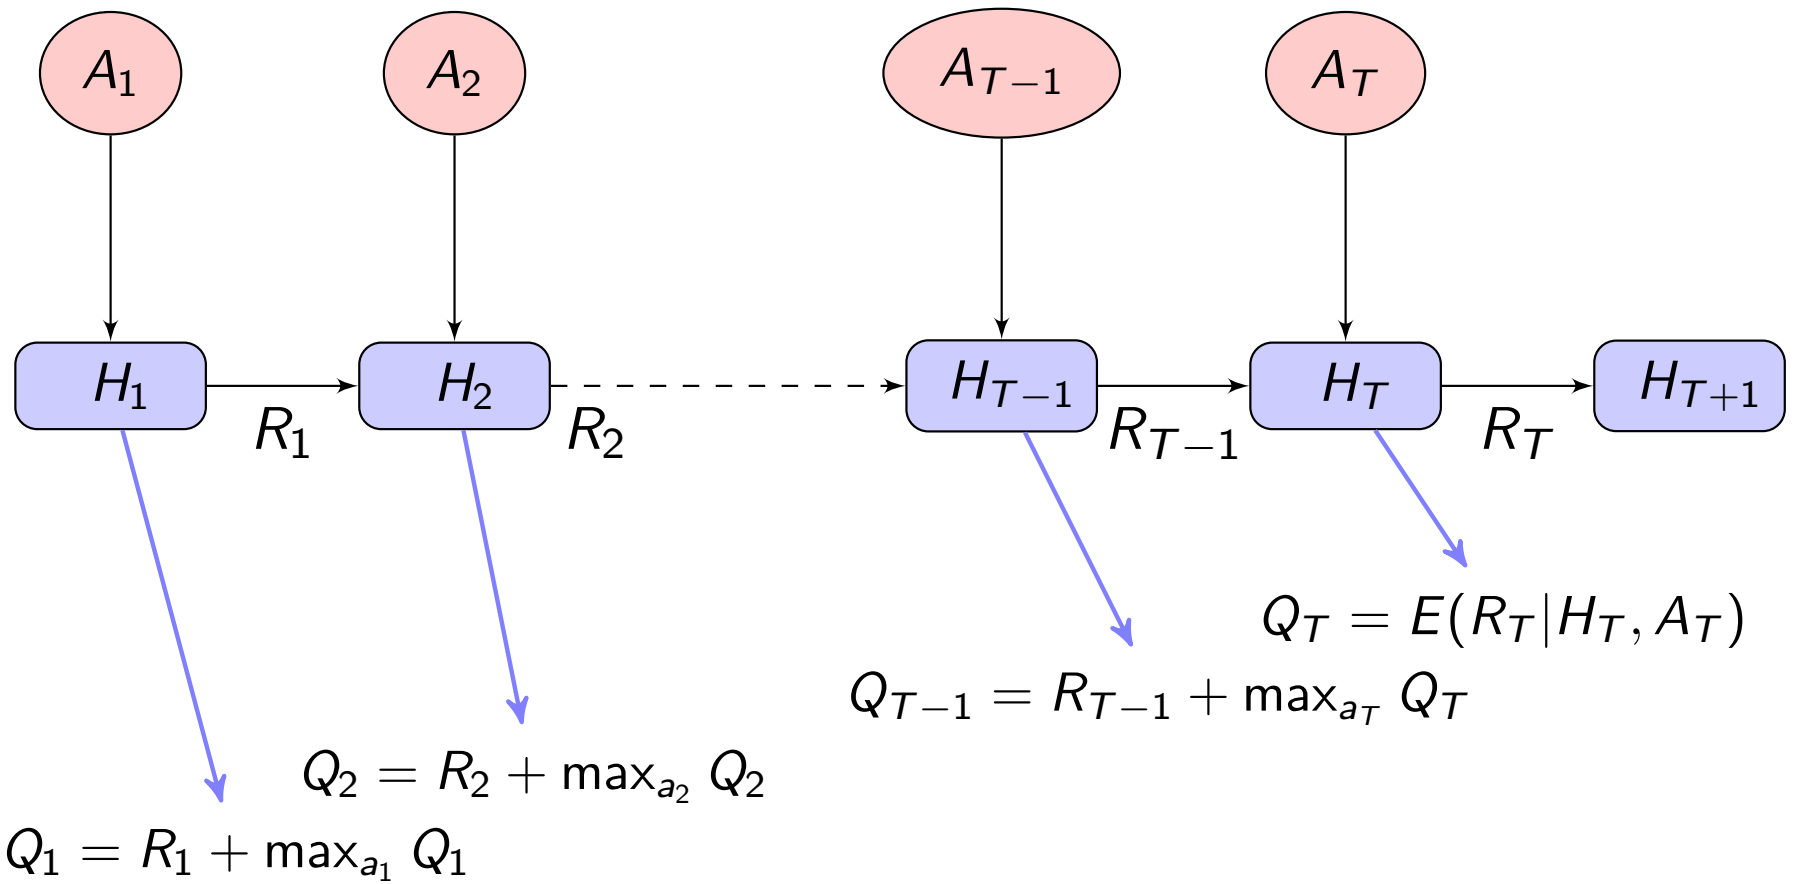

Construct a DTR via Q-learning

- An extension of regression to sequential treatments.
- Q learning with regression: estimate the Q-functions from data using regression and then find the optimal DTR.
- Decision not shared; rely on the assumption that models are correct.

Stabilized Dynamic Treatment Regimes

- Learn the optimal DTRs at all stages simultaneously.
- Identify an estimator of $\mathcal{V}(d)$, and directly maximize $\hat{\mathcal{V}}(d)$ over $d \in \mathcal{D}$, where $\mathcal{D}$ is a prespecified class of DTRs of interest, e.g., linear DTRs
- An inverse probability of treatment estimator (IPWE) of $\mathcal{V}(d)$

$$\hat{V}^{IPWE}(d) = \mathbb{P}_n \left(\frac{(\sum_j R_j) I\{A_j = d_j(H_j), j = 1, \ldots, T\}}{\prod_{j=1}^{T} P(A_j = d_j(H_j)|H_j)} \right).$$

Stabilized Dynamic Treatment Regimes

- Let $d_t(h_t) = I(h_t^T \beta > 0)$ and find $\hat{\beta}$ that maximizes $\hat{V}^{IPWE}(d)$.
- Replace a concave surrogate for the indicator to alleviate the computation difficulties.
- $P(A_t|H_t)$ can be estimated using e.g., logistic regression
- Apply the LASSO penalty for sparsity.

Simulation Studies: Generative Model

- Baseline covariates $X_{1,1}, \ldots, X_{1,50} \sim N(0,1)$
- $X_{t,1} = X_{t-1,1} + N(0,1), j \geq 2.$
- $\text{logit}(P(A_t = 1|H_t)) = 2X_t - 1$
- $Y \sim 10 + X_{1,1}X_{1,2} + X_{1,3} - \sum_{t=1}^{4} |2X_{t,1} - 3|\{I(A_t > 0) - I(X_{t,1}^3 - 0.5 > 0)\}^2 + N(0,1).$
- Optimal decisions: $d_t^*(h_t) = \text{sign}(X_{t,1} - 0.79)$

Simulation Studies

- Training data sample size 1000.
- Testing data sample size 10000.
- 500 replications.
- Evaluate using the values of the estimated DTRs.

Simulation Studies

	$n = 1000$
SDTR	8.66
Q learning	5.47
Shared Q learning	7.81

Medication regimes

- Large number of combinations of therapies are possible
- Use pharmacy claims to identify individual medications
- Group medications into therapeutic classes and identify the most frequent drug regimens

Missing data

- A significant portion of our data set is missing.

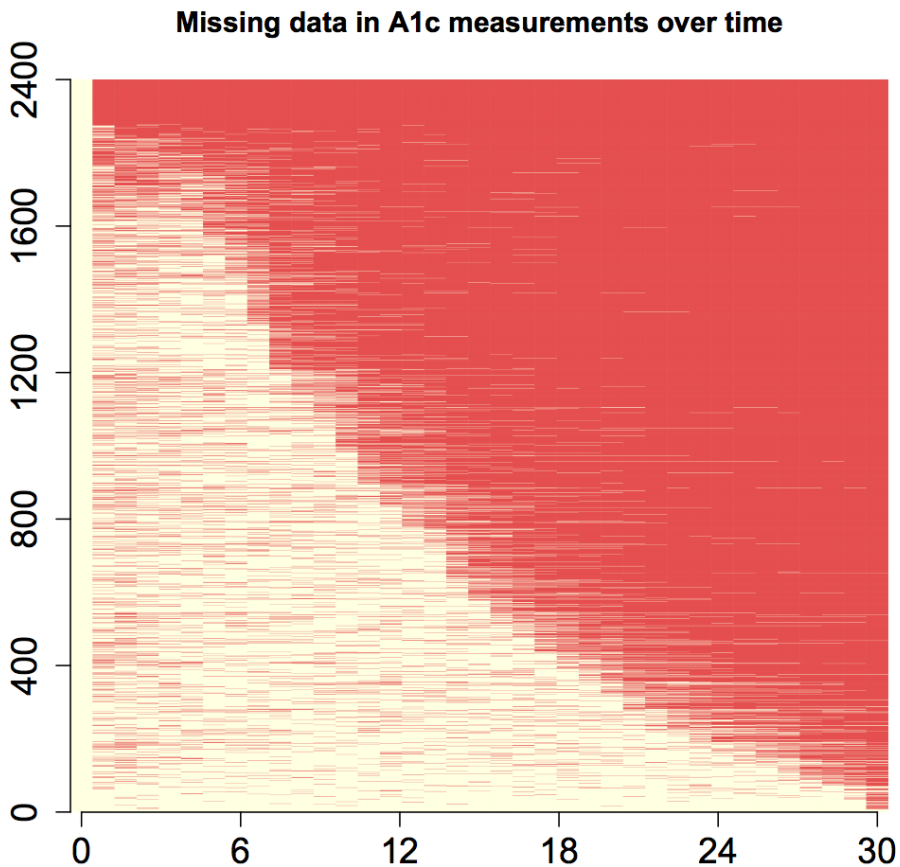

Missing data

- A significant portion of our data set is missing.

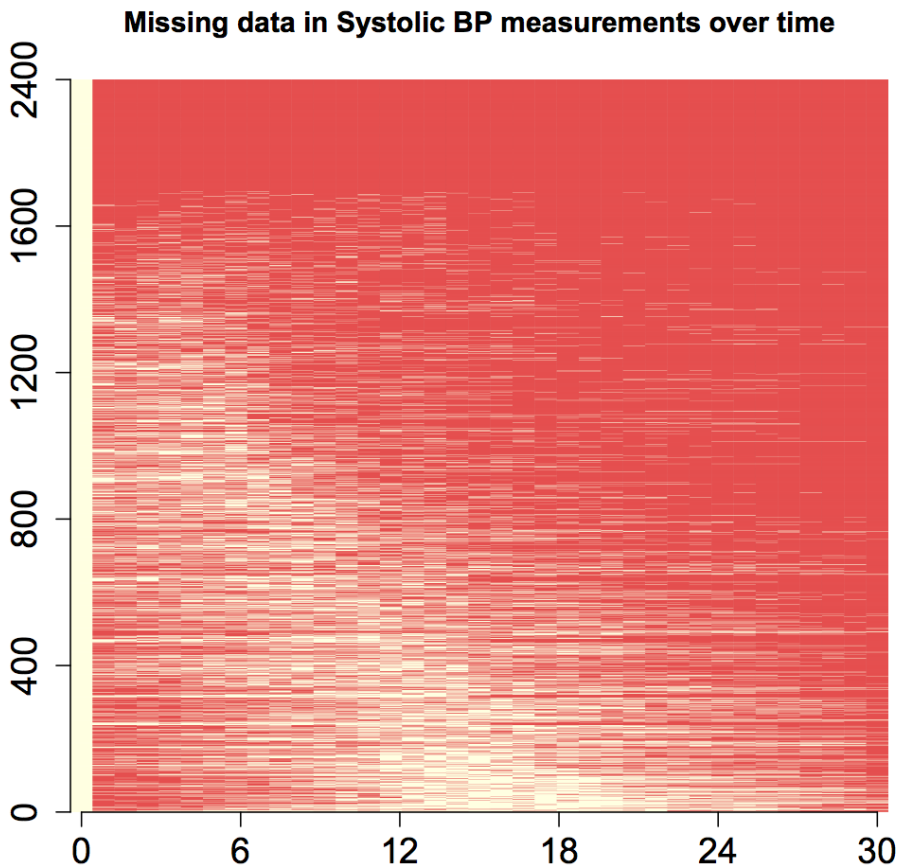

Missing Data

- EHR data are collected in a non-prescheduled fashion only when the patient seeks care or the physician orders care, creating a situation with intermittent missing data.
- The visiting process could be potentially informative about the patients' risk categories; not missing at random.

Multiple imputation

- The missing data are filled in m times to generate m complete data sets
- Estimate the optimal stabilized DTR using each complete data set.
- Results from the m complete data sets are combined.
- Fully Conditional Specification/Bayesian Mixed Effects Method to impute (Shortreed et al, 2014)
- Explore Bayesian multiple imputation approach and pattern mixture model methods.

Thanks!