Introduction to Personalized Medicine and Dynamic Treatment Regimes

Marie Davidian

Department of Statistics
North Carolina State University

http://www4.stat.ncsu.edu/~davidian

What is personalized medicine?

What is personalized medicine?

Personalized Medicine

Source of graphic: http://www.personalizedmedicine.com/

Personalized medicine

One size does not fit all:

- Multiple *treatment options*
- Patient *heterogeneity*

Basic premise: A *patient's characteristics* are implicated in which *treatment option* s/he should receive

- Genetic/genomic, other omic,...
- Demographic, physiologic/clinical measures, medical history,...

Popular perspective on personalized medicine

Subgroup identification/targeted treatment:

- Are there subgroups of patients who are *more likely* to do better on one treatment than on another?
- Can *biomarkers* be developed to identify such patients?
- Can a treatment be developed that *targets* a subgroup that is very likely to benefit from that treatment?

Popular perspective on personalized medicine

Subgroup identification/targeted treatment:

- Are there subgroups of patients who are *more likely* to do better on one treatment than on another?
- Can *biomarkers* be developed to identify such patients?
- Can a treatment be developed that *targets* a subgroup that is very likely to benefit from that treatment?

Focus: Targeting treatment to *subgroups* of the population

- "*Right patient for the treatment*"

Another perspective on personalized medicine

Clinical practice: Clinicians make (a series of) *treatment decision*(*s*) over the course of a patient's disease or disorder

- Key *decision points* in the disease process
- *Fixed schedule*, *milestone* in the disease process, *event* necessitating a decision
- Multiple *treatment options* at each
- *Accrued information* on the patient

Another perspective on personalized medicine

Clinical practice: Clinicians make (a series of) *treatment decision*(*s*) over the course of a patient's disease or disorder

- Key *decision points* in the disease process
- *Fixed schedule*, *milestone* in the disease process, *event* necessitating a decision
- Multiple *treatment options* at each
- *Accrued information* on the patient

Focus: *Given information on patient's characteristics*, can we determine the treatment(s) from among the available options most likely to benefit him/her?

- "*Right treatment for the patient*"

Another perspective on personalized medicine

Clinical practice: Clinicians make (a series of) *treatment decision*(*s*) over the course of a patient's disease or disorder

- Key *decision points* in the disease process
- *Fixed schedule*, *milestone* in the disease process, *event* necessitating a decision
- Multiple *treatment options* at each
- *Accrued information* on the patient

Focus: *Given information on patient's characteristics*, can we determine the treatment(s) from among the available options most likely to benefit him/her?

- "*Right treatment for the patient*"

This is the perspective we will take

Clinical decision-making

How are these decisions made?

- *Clinical judgment*
- *Practice guidelines* based on combining *study results* and *expert opinion*
- Synthesize all *information* on a patient up to the point of the decision to determine the *next treatment action*

How can statistics and statistical thinking be used to formalize clinical decision-making?

Dynamic treatment regime

Formalizing personalized medicine: At any *decision*

- Need a *rule* that takes as *input* the *accrued information* on the patient to that point and *outputs* the next treatment from among the *possible options*

Dynamic treatment regime: A set of formal such *rules*, each corresponding to a *decision point*

- *Dynamic* because the treatment action *varies* depending on the *accrued information*

Single decision

Simple example: Which treatment to give to patients who present with primary operable *breast cancer*?

- *Treatment options*: L-phenylalanine mustard and 5-fluorouracil (c_1) or c_1 + tamoxifen (c_2)
- *Information*: age, progesterone receptor level (PR)

Example rule: "If age $<$ 50 years and PR $<$ 10 fmol, give c_1 (coded as 1); otherwise, give c_2 (coded as 0)"

- *Mathematically*, the rule d is (*rectangular region*)

$$d(\text{age}, \text{PR}) = I(\text{age} < 50 \text{ and PR} < 10)$$

- *Alternatively*: Rules of form (*hyperplane*)

$$d(\text{age}, \text{PR}) = I\{\text{age} + 8.7\log(\text{PR}) - 60 > 0\}$$

Multiple decision points

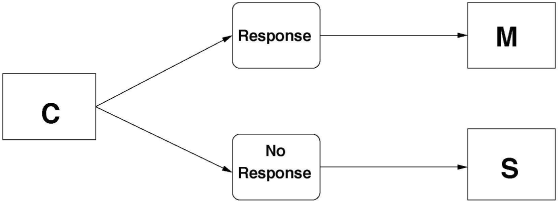

Two decision points:

- *Decision 1* : Induction chemotherapy (options c_1, c_2)
- *Decision 2* :
  - Maintenance treatment for patients who *respond* (options m_1, m_2)
  - Salvage chemotherapy for those who *don't* (options s_1, s_2)

Multiple decision points

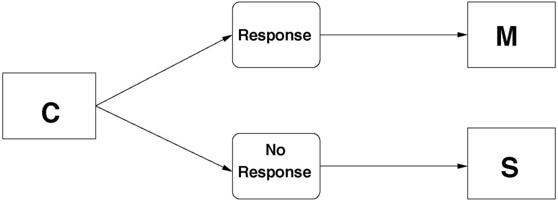

- *At baseline:* Information x_1 (*accrued information* $h_1 = x_1$)
- *Decision 1:* Two *options* $\{c_1, c_2\}$; *rule 1:* $d_1(x_1) \Rightarrow x_1 \rightarrow \{c_1, c_2\}$
- *Between decisions 1 and 2:* Collect *additional information* x_2, including *responder status*
- *Accrued information* $h_2 = \{x_1$, chemotherapy at decision 1, $x_2\}$
- *Decision 2:* Four *options* $\{m_1, m_2, s_1, s_2\}$; *rule 2:* $d_2(h_2) \Rightarrow h_2 \rightarrow \{m_1, m_2\}$ (responder), $h_2 \rightarrow \{s_1, s_2\}$ (nonresponder)
- *Regime*: $d = (d_1, d_2)$

Summary

Single decision: 1 decision point

- *Baseline information* $x \in \mathcal{X}$
- Set of *treatment options* $a \in \mathcal{A}$
- *Decision rule* $d(x)$, $d : \mathcal{X} \rightarrow \mathcal{A}$
- *Treatment regime:* d

Multiple decisions: K decision points

- *Baseline information* x_1, *intermediate information* x_k between decisions $k-1$ and k, $k = 2, \ldots, K$
- Set of *treatment options* at each decision k: $a_k \in \mathcal{A}_k$
- *Accrued information* $h_1 = x_1 \in \mathcal{H}_1$,

$$h_k = \{x_1, a_1, x_2, a_2, \ldots, x_{k-1}, a_{k-1}, x_k\} \in \mathcal{H}_k, \quad k = 2, \ldots, K$$

- *Decision rules* $d_1(h_1), d_2(h_2), \ldots, d_K(h_K)$, $d_k : \mathcal{H}_k \rightarrow \mathcal{A}_k$
- *Treatment regime* $d = (d_1, d_2, \ldots, d_K)$

“Best” treatment regime?

Obviously: There is an *infinitude* of possible regimes d

- $\mathcal{D}$ = class of *all possible* dynamic treatment regimes
- Can we find the “*best*” set of rules; i.e., the “*best*” dynamic treatment regime in $\mathcal{D}$?
- How to define “*best?*”

Approach: Formalize *best* using ideas from *causal inference*

Clinical outcome

Outcome: There is a *clinical outcome* by which treatment *benefit* can be assessed

- Survival time, CD4 count, indicator of no myocardial infarction within 30 days, ...
- *Larger outcomes* are *better*

Defining "best"

An optimal regime d^{opt}: Should satisfy

- If *all patients* in the *population* were to receive treatment according to d^{opt}, the *expected outcome* for the population would be *as large as possible*
- If an individual patient *were to receive treatment* according to d^{opt}, his/her *expected outcome* would be *as large as possible given the information available on him/her*

Formalize this. . .

Optimal regime for single decision

For simplicity: Consider regimes involving a *single decision* with *two* treatment options (coded as 0 and 1)

- $\mathcal{A} = \{0, 1\}$
- *Baseline information* $x \in \mathcal{X}$

Treatment regime: A single *rule* $d(x)$

- $d : \mathcal{X} \rightarrow \{0, 1\}$
- $d \in \mathcal{D}$, the class of *all* regimes

Potential outcomes

Treatments 0 and 1: We can hypothesize *potential outcomes*

- $Y^*(1)$ = outcome that *would be achieved* if a patient *were to receive* treatment 1
- $Y^*(0)$ similarly
- $E\{Y^*(1)\}$ is the *expected outcome* if *all patients* in the population were to receive 1; $E\{Y^*(0)\}$ similarly

In fact: In a *randomized trial* comparing 0 and 1, the *question of interest* can be stated as

> "Is the *expected outcome* if *all patients in the population* were to receive 0 *different from* that if they all were to receive 1 instead?"

- *Comparison* of $E\{Y^*(1)\}$ and $E\{Y^*(0)\}$

Potential outcome for a regime

Potential outcome for a regime: For any $d \in \mathcal{D}$, define $Y^*(d)$ to be the *potential outcome* for a patient with *baseline information* X if s/he were to receive treatment *in accordance with regime* d

$$Y^*(d) = Y^*(1)d(X) + Y^*(0)\{1 - d(X)\}$$

- *Recall:* d: $\mathcal{X} \rightarrow \{0, 1\}$
- *Potential outcome for regime d* for a patient with information X is the potential outcome for the treatment option *dictated by* d

Potential outcome for a regime

In general: For a *single decision* with set of treatment options $\mathcal{A}$ with elements a

- *Potential outcomes* $Y^*(a)$, $a \in \mathcal{A}$
- *Potential outcome* for regime $d \in \mathcal{D}$

$$Y^*(d) = \sum_{a \in \mathcal{A}} Y^*(a)\, I\{d(X) = a\}$$

Optimal dynamic treatment regime

Thus:

- $E\{Y^*(d)|X=x\}$ is the *expected outcome* for a patient with *information* x if s/he were to receive treatment according to regime $d \in \mathcal{D}$
- $E\{Y^*(d)\} = E[\,E\{Y^*(d)|X\}\,]$ is the expected (average) outcome for the *population* if *all patients* were to receive treatment according to regime $d \in \mathcal{D}$

Optimal regime: d^{opt} is a regime in $\mathcal{D}$ such that

- $E\{Y^*(d)\} \leq E\{Y^*(d^{opt})\}$ for all $d \in \mathcal{D}$
- $E\{Y^*(d)|X=x\} \leq E\{Y^*(d^{opt})|X=x\}$ for all $d \in \mathcal{D}$ and all $x \in \mathcal{X}$

Optimal dynamic treatment regime

Simplest case: $\mathcal{A} = \{0, 1\}$

$$Y^*(d) = Y^*(1)d(X) + Y^*(0)\{1 - d(X)\}$$

- *Thus*

$$\begin{aligned} E\{Y^*(d)\} &= E\Big[E\{Y^*(d)|X\}\Big] \\ &= E\Big[E\{Y^*(1)|X\}d(X) + E\{Y^*(0)|X\}\{1 - d(X)\}\Big] \end{aligned}$$

- It *follows* that

$$d^{opt}(x) = I\Big[E\{Y^*(1)|X = x\} > E\{Y^*(0)|X = x\}\Big]$$

Important philosophical point

Distinguish between:

- The "*best*" treatment for a patient
- The "*best*" treatment *decision* for a patient *given the information available on the patient*

Best treatment for a patient: Option $a^{best} \in \mathcal{A}$ corresponding to the *largest* $Y^*(a)$ for the patient

Best treatment given the information available:

- We *cannot* hope to determine a^{best} because we can never see *all* potential outcomes on a given patient
- We *can* hope to make the *optimal decision* given the *information available*, i.e., find d^{opt} that makes $E\{Y^*(d)\}$ and $E\{Y^*(d)|X = x\}$ *as large as possible*

Statistical problem

Goal: Given *data* from a clinical trial or observational study, *estimate* an *optimal regime* d^{opt} satisfying this definition

Observed data: Single decision, $K = 1$, $\mathcal{A} = \{0, 1\}$.

$$(X_i, A_i, Y_i), \quad i = 1, \ldots, n \quad \text{iid}$$

- X = *baseline* information
- A = treatment *actually received*
- Y = *observed* outcome

Required:

- *Optimal regime* is defined in terms of *potential outcomes*
- Must *express equivalently* in terms of *observed data*

Statistical problem

Consistency assumption:

$$Y = Y^*(1)A + Y^*(0)(1 - A)$$

No unmeasured confounders assumption:

$$Y^*(0), Y^*(1) \perp\!\!\!\perp A_1 \mid X_1$$

- Automatically satisfied in a *clinical trial*
- *Unverifiable* but necessary in an *observational study*

Statistical problem

Consistency assumption:

$$Y = Y^*(1)A + Y^*(0)(1-A)$$

No unmeasured confounders assumption:

$$Y^*(0), Y^*(1) \perp\!\!\!\perp A_1 \mid X_1$$

- Automatically satisfied in a *clinical trial*
- *Unverifiable* but necessary in an *observational study*

Under these assumptions:

$$E\{Y^*(d)\} = E\Big[E(Y|X_1, A=1)d(X_1) + E(Y|X_1, A=0)\{1-d(X_1)\}\Big]$$

- *Butch* will show how this leads to an estimator for d^{opt}

Multiple decision points

Potential outcomes: Same ideas, only more *complicated*

- *Baseline information* X_1
- K decisions, set of *treatment options* $\mathcal{A}_k$, $k = 1, \ldots, K$
- $\bar{a}_k = (a_1, \ldots, a_k)$ = possible *treatment history* through decision k, taking values in $\bar{\mathcal{A}}_k = \mathcal{A}_1 \times \cdots \mathcal{A}_k$
- $X_k^*(\bar{a}_{k-1})$ = intermediate information between decisions $k - 1$ and k that *would arise* for a patient if s/he *were to receive* treatment history $\bar{a}_{k-1}$
- $Y^*(\bar{a}_K)$ = outcome that *would be achieved* for a patient if s/he *were to receive* treatment history $\bar{a}_K$
- *Potential outcomes* under treatment history $\bar{a}_K$

$$X_2^*(a_1), X_3^*(\bar{a}_2), \ldots, X_K^*(\bar{a}_{K-1}), Y^*(\bar{a}_K)$$

Multiple decision points

Potential outcomes for regime $d = (d_1, \dots, d_K)$:

- Write $\bar{d}_k = (d_1, \dots, d_k)$ = *first k rules* in d, $\bar{d}_K = d$
- Write $\bar{x}_k$ similarly
- $h_k = (x_1, a_1, x_2, \dots, a_{k-1}, x_k) = (\bar{x}_k, \bar{a}_{k-1})$
- With $X_1 = x_1 = h_1$

$$d_1(h_1) = a_1, X_2^*(d_1) = X_2^*(a_1) = x_2, d_2(h_2) = a_2,$$

$$X_3^*(\bar{d}_2) = X_3^*(\bar{a}_2) = x_3, \dots, d_{K-1}(h_{K-1}) = a_{K-1},$$

$$X_K^*(\bar{d}_{K-1}) = X_K^*(\bar{a}_{K-1}) = x_K, d_K(h_K) = a_K, Y^*(d) = Y^*(\bar{a}_K) = y$$

Multiple decision points

Potential outcomes for regime $d = (d_1, \dots, d_K)$:

- Write $\bar{d}_k = (d_1, \dots, d_k)$ = *first k rules* in d, $\bar{d}_K = d$
- Write $\bar{x}_k$ similarly
- $h_k = (x_1, a_1, x_2, \dots, a_{k-1}, x_k) = (\bar{x}_k, \bar{a}_{k-1})$
- With $X_1 = x_1 = h_1$

$$d_1(h_1) = a_1, X_2^*(d_1) = X_2^*(a_1) = x_2, d_2(h_2) = a_2,$$

$$X_3^*(\bar{d}_2) = X_3^*(\bar{a}_2) = x_3, \dots, d_{K-1}(h_{K-1}) = a_{K-1},$$

$$X_K^*(\bar{d}_{K-1}) = X_K^*(\bar{a}_{K-1}) = x_K, d_K(h_K) = a_K, Y^*(d) = Y^*(\bar{a}_K) = y$$

- $X_k^*(\bar{d}_{k-1})$ is the information that *would arise* between decisions $k-1$ and k if a patient *were to receive* the treatments dictated by the first $k-1$ rules in d
- $Y^*(d)$ is the outcome a patient *would achieve* if s/he *were to receive* the K treatments according to the K rules in d

Optimal dynamic treatment regime

Thus: For $K \geq 1$ decision points

- $E\{Y^*(d)|X_1 = x_1\}$ is the *expected outcome* for a patient with *baseline information* x_1 if s/he were to receive all K subsequent treatments according to regime $d \in \mathcal{D}$
- $E\{Y^*(d)\} = E[\,E\{Y^*(d)|X\}\,]$ is the expected (average) outcome for the *population* if *all patients* were to receive all K treatments according to regime $d \in \mathcal{D}$

Optimal regime: d^{opt} is a regime in $\mathcal{D}$ such that

- $E\{Y^*(d)\} \leq E\{Y^*(d^{opt})\}$ for all $d \in \mathcal{D}$
- $E\{Y^*(d)|X_1 = x_1\} \leq E\{Y^*(d^{opt})|X_1 = x_1\}$ for all $d \in \mathcal{D}$ and all $x_1 \in \mathcal{X}_1 = \mathcal{H}_1$

Statistical problem

Goal: Given *data* from a clinical trial or observational study, *estimate* an *optimal regime* d^{opt} satisfying this definition

Observed data: K decisions

$$(X_{1i}, A_{1i}, X_{2i}, A_{2i}, \dots, X_{(K-1)i}, A_{(K-1)i}, X_{Ki}, A_{Ki}, Y_i), \quad i = 1, \dots, n, \quad iid$$

- Y = *observed* outcome
- $A_k, k = 1, \dots, K$ = treatment *received* at decision k
- $X_k, k = 2, \dots, K$ = *intermediate information* observed between decisions $k-1$ and k
- $H_k = (X_1, A_1, X_2, \dots, A_{k-1}, X_k), k = 2, \dots, K$ = *accrued information* to decision k

Consistency: $X_k = X_k^*(\bar{A}_{k-1}), k = 2, \dots, K; \ Y = Y^*(\bar{A}_K)$

Statistical problem

Critical assumption: *Sequential randomization*

- Generalization of *no unmeasured confounders*
- Treatment received at *each decision* $k = 1, \dots, K$ is *statistically independent* of potential outcomes that would be achieved under all treatment options *at all decisions* given the *accrued information* H_k

$$W^* = \{X_2^*(a_1), X_3^*(\bar{a}_2), \dots, X_k^*(\bar{a}_{k-1}), \dots, X_K^*(\bar{a}_{K-1}), Y^*(\bar{a}_K) \text{ for all } \bar{a}_K \in \bar{\mathcal{A}}_K\}$$

$$W^* \perp\!\!\!\perp A_k | H_k$$

- *Butch* will characterize the form of d^{opt} and show how this leads to an expression for d^{opt} in terms of *observed data* and *estimators* for d^{opt}

Data sources

Studies:

- Longitudinal observational
- Sequential, Multiple Assignment, Randomized Trial

Cancer example: Randomization at •s

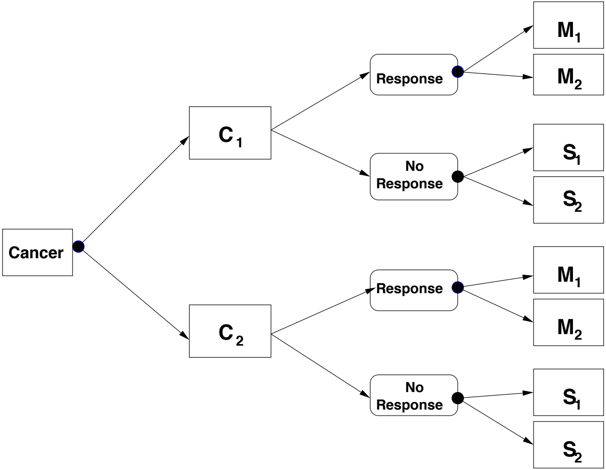

Sequential randomization automatically holds

Summary

- Formalizing this problem in terms of *potential outcomes* yields *precise definition* of an optimal regime
- And *clarifies* what must be assumed about *observed data* for them to be used to *estimate* an optimal regime

Reference

Schulte, P. J., Tsiatis, A. A., Laber, E. B., and Davidian, M. (2014). Q- and A-learning methods for estimating optimal dynamic treatment regimes. *Statistical Science*, in press.