Bayesian Nonparametric Statistics: A New Toolkit for Discovery in Cancer Research

Peter F. Thall
Department of Biostatistics
The University of Texas
MD Anderson Cancer Center

PSI Journal Club Webinar
July 2019

Collaborators: Peter Mueller, Yanxun Xu, Michele Guindani, William Hua, Abdus Wahed, Reza Mehran, Borje Andersson

Outline

1. Bayesian Nonparametric (BNP) Models
2. BNP-Model-Based Clinical Trial design: Comparing Treatments to Control Lung Air Leaks After Lung Surgery
3. Personalized Medicine: Optimizing PK-Guided Dose of IV Busulfan in Allogeneic Stem Cell Transplantation, Based on Disease Status and Age
4. Correcting for Bias: Estimating Mean Survival Times of 16 Dynamic Treatment Regimes for Acute Leukemia

In the Beginning . . . Thomas Ferguson (1973) proposed the Dirichlet process (DP) prior for a random probability distribution F:

1. A DP(α_0, G_0) prior has scale parameter $\alpha_0 > 0$ and base probability measure G_0, represented using the "stick-breaking" construction (Sethuraman, 1974) as $G = \sum_{h=1}^{\infty} w_h \delta_{\theta_h}$, where $\theta_h \sim iid\ G_0$, and δ_{θ_h} is the Dirac delta function with mass 1 at θ_h, and the random weights are

$$w_h = v_h \prod_{l<h} (1 - v_l) \quad \text{with} \quad v_h \sim \text{beta}(1, \alpha)$$

2. Under a DP(α_0, G_0) prior on F, for any partition $\{A_1, \cdots, A_r\}$ of the domain of F, the probability vector

$$[F(A_1), \cdots, F(A_r)] \sim Dirichlet[\alpha_0 G_0(A_1), \cdots, \alpha_0 G_0(A_r)].$$

3. Given sample $Y_1, \cdots, Y_n \sim$ *iid* F, the posterior distribution is

$$(F \mid Y_1, \cdots, Y_n) \sim DP\left(\alpha_0 + n, \frac{\alpha_0 G_0 + \sum_{i=1}^{n} \delta_{Y_i}}{\alpha_0 + n}\right).$$

4. Since any DP puts all of its probability mass on discrete distributions F, to extend this to include continuous distributions, one may assume a DP mixture model by replacing δ_{θ_h} with a continuous kernel, most often a normal $N(\theta_h, \sigma^2)$, giving

$$G = \sum_{h=1}^{\infty} w_h \, N(\theta_h, \sigma^2).$$

5. More generally, $\{\theta_1, \theta_2, \cdots\}$ in $G = \sum_{h=1}^{\infty} w_h \delta_{\theta_h}$ may be *any random objects*: Real numbers, matrices, images, graphs, dogs, cats, etc. Most BNP priors put probability into any neighborhood of any model (they have "full support") $\Longrightarrow$ Many BNP models reveal patterns in data that conventional models miss.

6. Steve MacEachern (1999) included regression on covariates $\boldsymbol{Z}$ by replacing each mean θ_h with a function $\theta_h(\boldsymbol{Z})$, to obtain the Dependent Dirichlet process (DDP). This gives F as a mixture of normals, and includes a stochastic process prior for $\{\theta_h(\boldsymbol{Z})\}$:

$$F(y \mid \boldsymbol{Z}) = \sum_{h=1}^{\infty} w_h \; N(y; \theta_h(\boldsymbol{Z}), \sigma^2).$$

7. A Gaussian process (GP) prior for a stochastic process $\{\theta(\boldsymbol{Z})\}$ is characterized by the marginal of any n-tuple $[\theta(\boldsymbol{Z}_1), \cdots, \theta(\boldsymbol{Z}_n)]$ having mean vector $[\mu(\boldsymbol{Z}_1), \cdots, \mu(\boldsymbol{Z}_n)]$ and $n \times n$ variance - covariance matrix $C(\boldsymbol{Z}_i, \boldsymbol{Z}_j)$ for any n and $\{\boldsymbol{Z}_1, \cdots, \boldsymbol{Z}_n\}$.

8. Assuming each $\theta_h(\boldsymbol{Z}) \sim GP(\mu_h, C)$ with $\mu_h(\boldsymbol{Z}) = \boldsymbol{\beta}_h \boldsymbol{Z} = \sum_{j=1}^{p} \beta_{h,j} Z_j$ with priors $\boldsymbol{\beta}_h \sim iid \; MN_p(\boldsymbol{\beta}_0, \Sigma_0)$ gives a Dirichlet Process with a Gaussian Process Prior (DDP-GP).

A Trial to Compare Treatments for Controlling Air Leaks After Lung Surgery

The Problem: Design a randomized trial to compare the gel sealant Progel to standard care (sutures + staples) to control intraoperative air leaks (IALs) after thoracotomies (lung resections).

Outcome: T = Days to resolve IAL, with T=0 if no air leak occurs

Background: Based on historical data, **mean time to resolve an IAL = 8 days**. A one-sided, two-sample, .05-level t-test to detect a 25% drop in the mean, from 8 to 6 days, with power = .80, would require n=476 patients. If we use Y = log(T+1) as the outcome and repeat the computation, n=280. $\longrightarrow$ **The trial is not feasible**

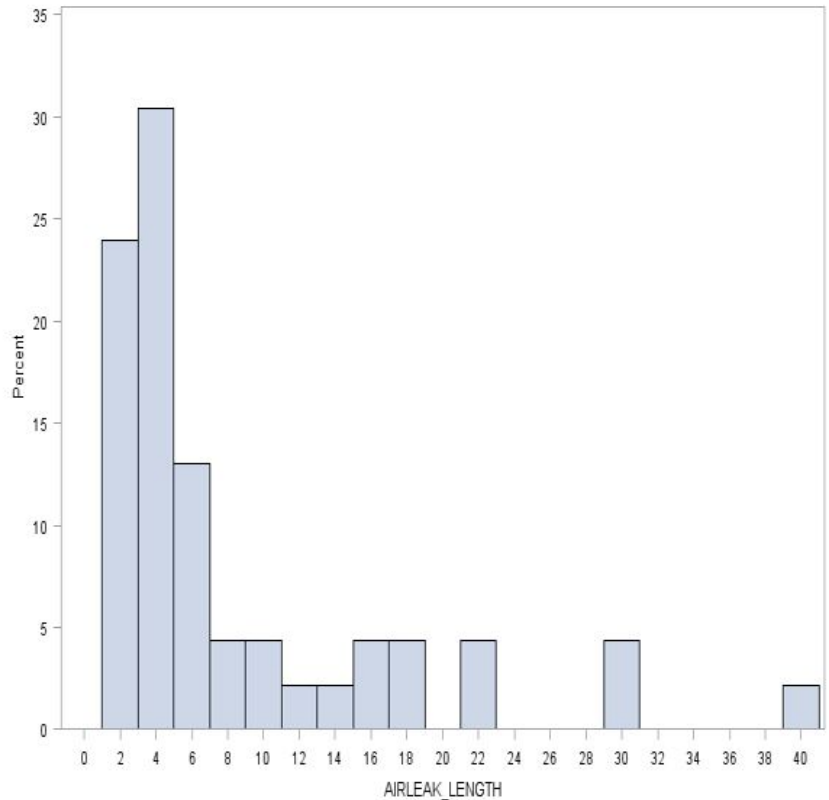

Historical distribution of T = days to resolve an IAL with Sutures + Staples is multi-model, and not symmetric bell shaped $\Longrightarrow$ A two-sample t-test is INCORRECT for constructing a trial design.

Solution: Apply Bayesian Nonparametric Modeling + Utility-Based Decision Criteria

A BNP-model-based design for a randomized trial to compare Progel sealant to standard use of staples and sutures to control intra-operative air leaks in patients undergoing lung resection.

The design is based on a Bayesian nonparametric model for the distribution of T that

- Accounts for the skewed shape of the distribution
- Accounts for multimodality (humps at large times)
- Allows the possibility that T=0 (No IAL develops)

Bayesian Nonparametric Model

Here, $J = 2$ treatments. Index $j = 0$ for Control, $j = 1$ for Progel, and let δ_0 denote the probability distribution with mass 1 at $T = 0$ (no IAL occurs). The assumed BNP distribution of $T \mid j$ is

$$\begin{aligned} G_j &= \nu_{j,0}\delta_0 + (1-\nu_{j,0})\sum_{h=1}^{\infty} w_h N(\theta_{j,h}, \sigma^2) \\ &= \nu_{j,0}\delta_0 + (1-\nu_{j,0})M_j \end{aligned}$$

where $\sum_{h=1}^{\infty} w_h = 1$. We impose $M_1 \prec M_0$, that is, Progel can only make things better, since it is inert and does not react with cells.

For each h, we assume $\theta_{0,h}, \theta_{1,h} \sim \textit{iid}\ M^* =$ a multivariate normal truncated base measure, allowing ties $\theta_{0,h} = \theta_{1,h}$ since time is in days, with most times $T \leq 20$ days.

Histogram of Historical (Control) IAL Resolution Times

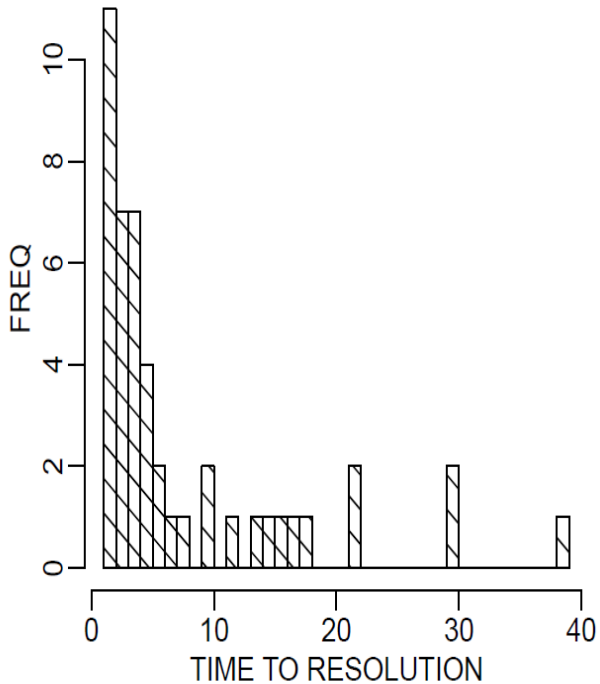

Bayesian Nonparametric Probability Density Estimates

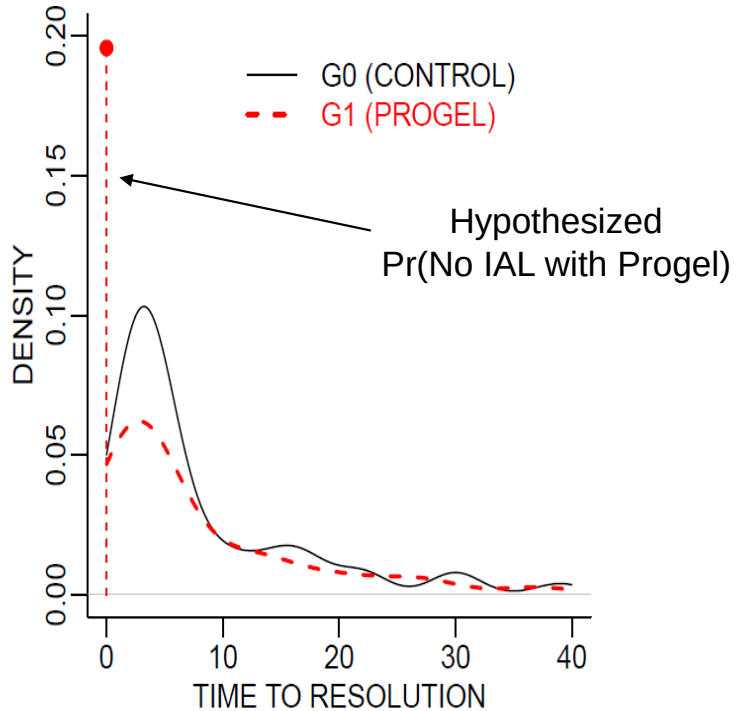

Elicited utilities of
T = days to resolve Intraoperative Air Leak

T (days)	0	5	10	15	20	25	30	35	≥ 40
Utility	100	50	10	6	5	4	3	2	0

The elicited utilities show that using T as the primary outcome would be very misleading.

Instead of T, we use its utility U(T)

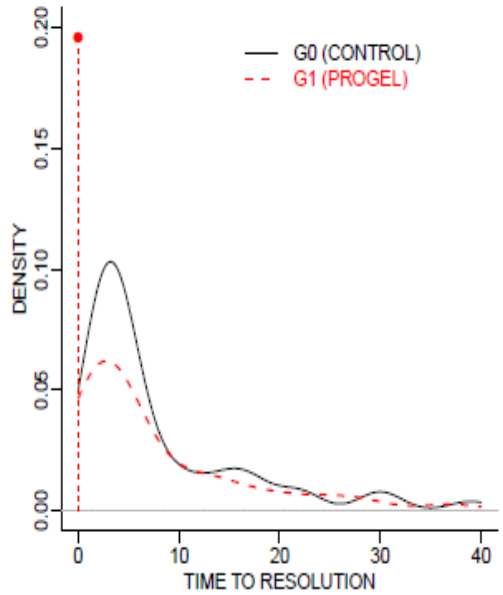

Bayesian Nonparametric Probability Density Estimates

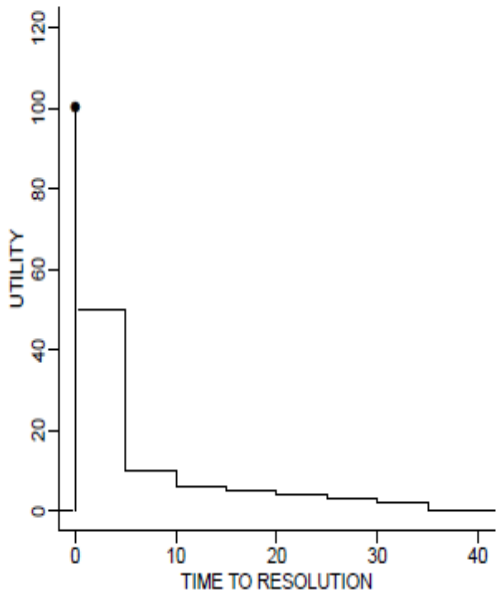

Elicited Utilities

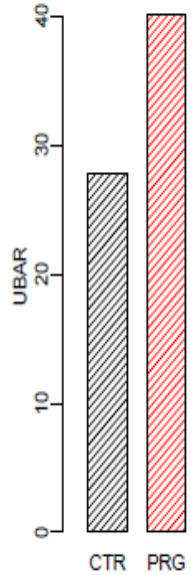

Mean Utilities

For mean utilities $\overline{U}_j = \int U(y) dG_j(y),\ j = 01$, under the BNP model, given interim or final data

$$\mathbf{Y}_n = \{Y_{ji}\, i = 1, \cdots, n/2, \quad j = 0, 1\}$$

the Decision Criterion is

$$\eta(\epsilon_U, \boldsymbol{y}_n) = P(\overline{U}_1 > \overline{U}_0 + \epsilon_U \mid \boldsymbol{Y}_n)$$

given targeted improvement $\epsilon_U \geq 0$, and

the Decision Rules are

- If $\eta(\epsilon_U, \boldsymbol{y}_n) > \overline{c}$ stop and conclude that Progel is superior.
- If If $\eta(\epsilon_U, \boldsymbol{y}_n) < \underline{c}$ stop for futility and conclude that Progel is not better than sutures and staples.

A Feasible Design:

Based on what the PI considered a meaningful improvement on the time domain, we used $\epsilon_U = 18$ utility points as the targeted improvement on the mean utility domain

Sample Size N = 48 patients, randomized fairly, enforcing exact balance between arms at n =16, 32, and 48

One-sided group sequential tests for superiority and futility, based on posterior decision criteria, are done at 16, 32, and 48 patients

Computer Simulations: Design Operating Characteristics

MSS = Mean Sample Size

TIE = Type I Error Probability

PCD = Probability of a Correct Decision

Scenario	$U_1^o - U_0^o$	MSS	TIE	PCD
1	0	16.80	0.00	1.00
1a	0	16.32	0.00	1.00
2	0	28.80	0.02	0.98
2a	0	28.00	0.01	0.99
3	41.64	29.12	-	1.00
4	19.15	40.16	-	0.63
5	29.68	31.68	-	0.93
6	34.15	29.92	-	0.94
7	43.47	28.64	-	0.96
8	8.13	34.08	-	0.79
9	10.25	34.88	-	0.74

Personalizing PK-Guided Dosing in Allogeneic Stem Cell Transplantation

• Intravenous (IV) busulfan is a key component of the preparative regimen in allogeneic stem cell transplantation for acute leukemia.

• Systemic busulfan exposure, characterized by area under the plasma concentration curve (AUC), is strongly associated with clinical outcome :

- AUC too high $\Rightarrow$ Severe life-threatening toxicities.
- AUC too low $\Rightarrow$ High risks of disease recurrence and graft failure, both increasing the risk of death.

• Each patient's optimal AUC interval must be determined for therapeutic use by first administering a pre-preclinical dose to determine their busulfan pharmacokinetics (PK)

A smoothed martingale residual plot of log(survival time) as a function of IV busulfan AUC in allogeneic stem cell transplant patients (2002) implied that **there is a 'middle interval' for delivered dose (AUC) that minimizes the risk of death**

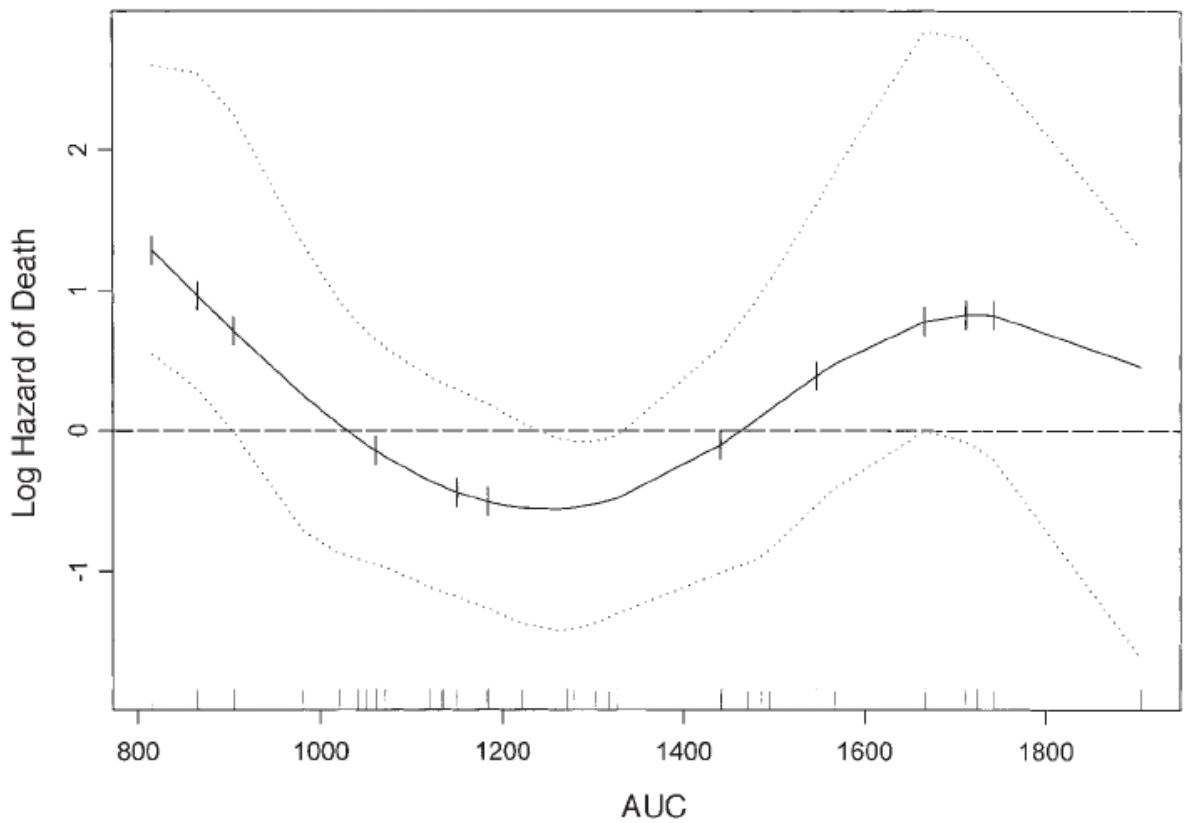

Getting Personal: Since Age and CR status at allosct (Yes = In Complete Remission, No = Active Disease) are strongly predictive of survival time in allosct: Can the optimal AUC interval be personalized to maximize expected survival time?

Historical Dataset: $n = 151$ allosct patients received a standard 4-day preparative regimen of fludarabine + IV busulfan. PK studies of IV busulfan were performed as an optional procedure, but the PK information was not used to adjust busulfan dose. $\Rightarrow$ **This historical dataset cannot be extended ethically**.

BNP Statistical Analysis: We fit a DDP-GP regression model for survival time with covariates $\boldsymbol{Z}$ = (AUC, Age, CR status).

Preliminary Estimates of the Survival Distribution

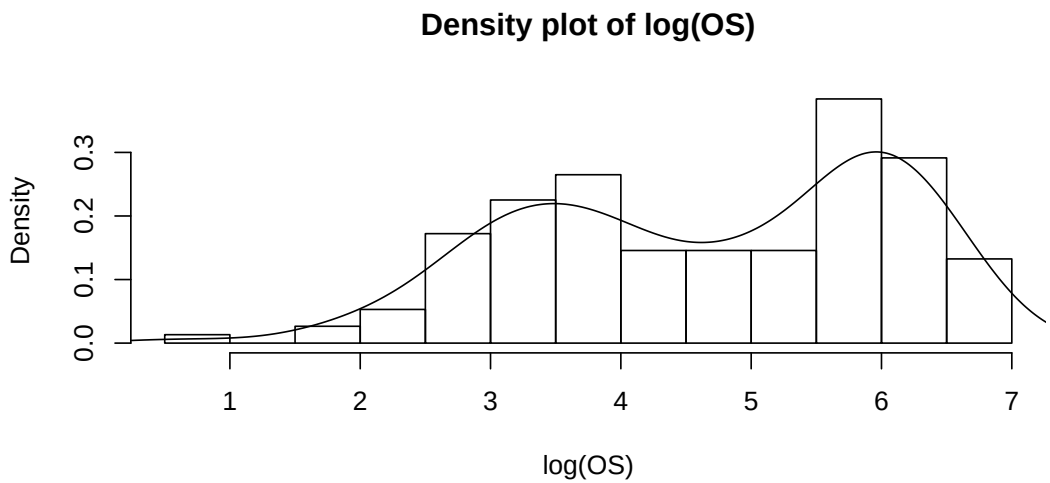

An Improved DDP-GP Survival Regression Model

Given the usual likelihood function

$$\mathcal{L}(\theta \mid \mathcal{D}_n) = \prod_{i=1}^{n} \{f_{\boldsymbol{Z}_i}(Y_i \mid \theta)\}^{\delta_i} \{1 - F_{\boldsymbol{Z}_i}(Y_i \mid \theta)\}^{1-\delta_i},$$

for general survival regression data $\mathcal{D}_n = \{Y_i, \delta_i, \boldsymbol{Z}_i\}_{i=1}^{n}$, Yanxun Xu extended the DDP-GP model of Xu et al. (2016) to have a more robust covariance matrix :

The regression of Y_i on $\boldsymbol{Z}_i$ is modeled as a Gaussian Process $\theta_h(\boldsymbol{Z}) \sim \text{GP}(\mu_h, C)$ with $\mu_h(\boldsymbol{Z}_i; \boldsymbol{\beta}_h) = \boldsymbol{Z}_i \boldsymbol{\beta}_h$ for $h = 1, 2, \cdots$, and

$$C(\boldsymbol{Z}_i, \boldsymbol{Z}_\ell) = \sigma_0^2 \ \exp\left\{ - \sum_{d=1}^{D} \frac{(Z_{id} - Z_{\ell d})^2}{\lambda_d^2} \right\} + \delta_{i\ell} J^2.$$

with jitter J^2, usually about .01, added to provide numerical stability by avoiding singular covariance matrices.

Posterior estimates of mean log(survival time) as a function of AUC, for each of eight (CR Status, Age) combinations, with 95% credible intervals, under the DDP-GP model

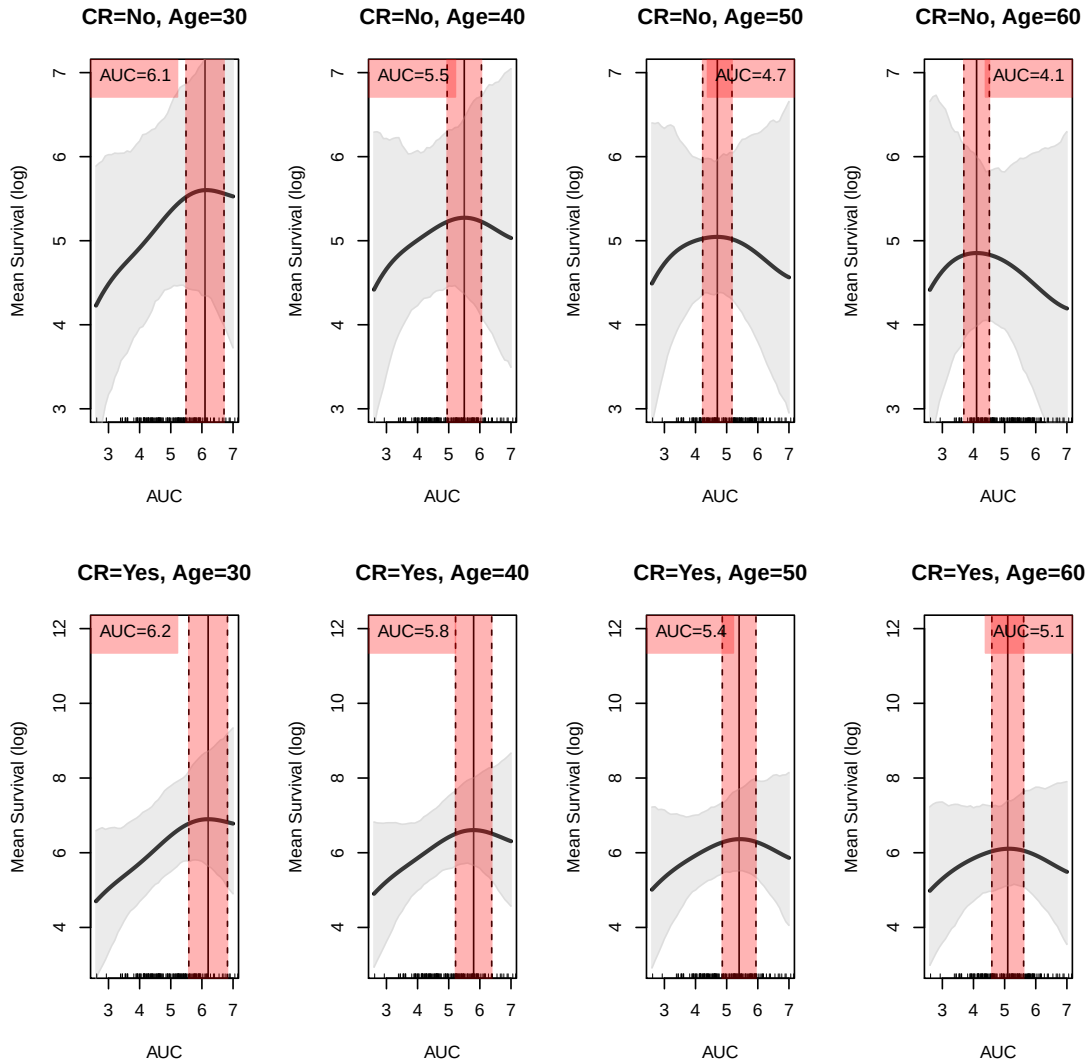

Estimated **Optimal Targeted Intervals** of IV Busulfan AUC
Personalized For Given (CR Status, Age)

We define the predicted optimal IV busulfan targeted AUC for future patient $n+1$ with covariates $\boldsymbol{Z}$ = (CR Status, Age, AUC) as

$$\widehat{AUC}_{n+1} = \text{argmax}_{AUC}\ E(Y_{n+1} \mid \boldsymbol{Z}, \mathcal{D}_n)$$

Since the laboratory error in evaluation of AUC is up to about 6%, the optimal AUC interval for future patient $n+1$ is defined as

$$[\ 0.9\ \widehat{AUC}_{n+1},\quad 1.1\ \widehat{AUC}_{n+1}\ ]$$

Estimated Optimal IV Busulfan AUC Intervals for (CR Status, Age)

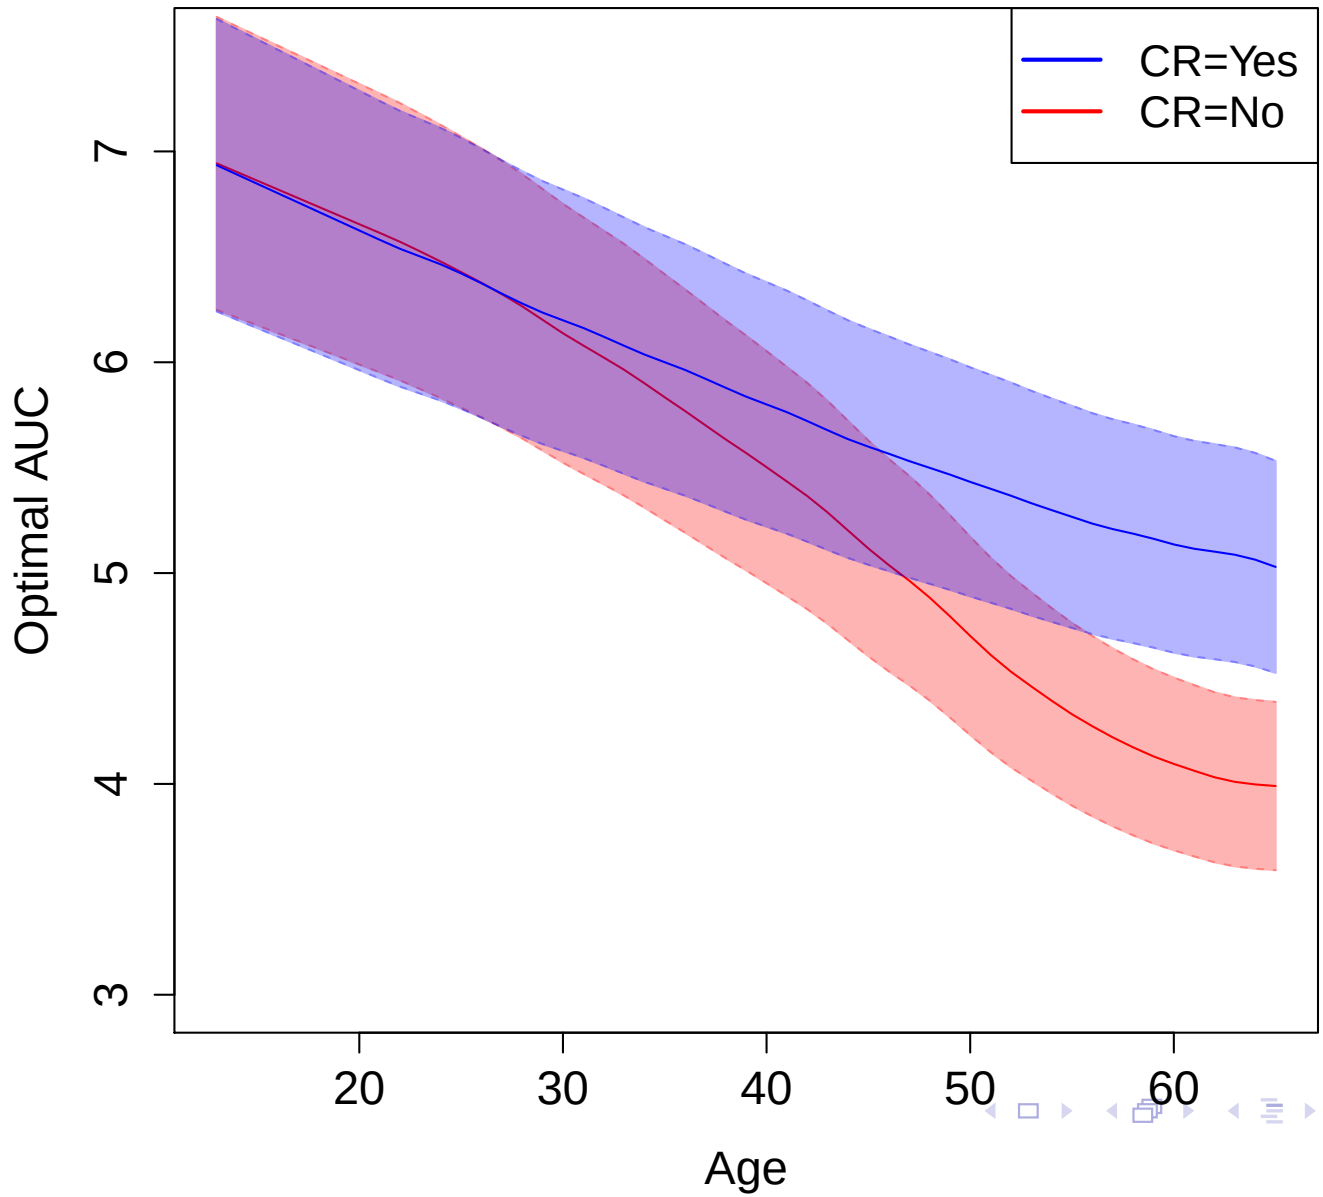

How to do your own Bayesian nonparametric survival analyses :

The DDP-GP model is a general tool for robust Bayesian survival regression analysis.

For implementation, the R package *DDPGPSurv* is available at

https://cran.r-project.org/web/packages/DDPGPSurv/index.html

Estimating Mean Survival Times with 16 Dynamic Treatment Regimes for Acute Leukemia

or

How We Discovered that a BNP Survival Regression Model Can Correct for Bias Better Than Frequentist Inverse Probability Weighting Methods

Acute Myelogenous Leukemia (AML) or Myelodysplastic Syndrome (MDS) :

- Characterized by high % of circulating blasts (leukemia cells) that never differentiate into functional blood cells
- Five-year survival is 35% - 50%, depending on prognostic covariates

Goal of Chemotherapy: Achieve CR = Complete Remission

1) Kill the leukemia cells,
2) Bring white blood cells back to functional level
3) Bring platelets back to functional level

Key Facts About Chemotherapy for AML/MDS

- Time to achieve CR is very important for predicting survival time (shorter time to CR is better)
- Disease may recur (progress) after CR is achieved
- Some patients may not achieve CR ('Resistant Disease')
- Salvage therapy is given if
  - Resistant Disease, or
  - at Progression after CR
- Patients may die
  - if disease is resistant
  - during induction chemo
  - while in CR
  - after progression

Schematic of the Trial Design

A Four-Arm Trial in AML/MDS (Estey, et al. *Blood*, 1999)
Motivated by success treating APL with ATRA

210 newly diagnosed poor prognosis AML/MDS patients were randomized among 4 induction treatment arms, balancing dynamically on patient prognostic covariates:

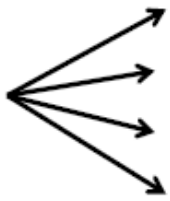

FAI (Fludarabine + ara-C + Idarubicin)
FAI + ATRA (all-trans retinoic acid, vitamin A)
FAI + G-CSF
FAI + ATRA + G-CSF

A 2x2 Factorial Design	
FAI	FAI + ATRA
FAI + G-CSF	FAI + ATRA + G-CSF

G-CSF Effect

ATRA Effect

Results of the 1999 Data Analysis

Based on **Kaplan-Meier plots, frequentist logistic regression for Pr(CR), and Cox model regression**, accounting for induction chemo arm and covariates :

No significant differences among the 4 chemo arms in terms of

- Pr(CR)
- Event-free survival time
- Overall survival time
- Survival time following CR

(all p-values = .18 to .99 for various tests)

→ No ATRA effect, No G-CSF effect

These analyses focused on the 4 frontline chemos only and ignored salvage therapy, as done routinely in the medical literature.

The 4 Potential Treatment-Outcome Pathways in Chemotherapy of AML/MDS

Keeping Track of Survival Time

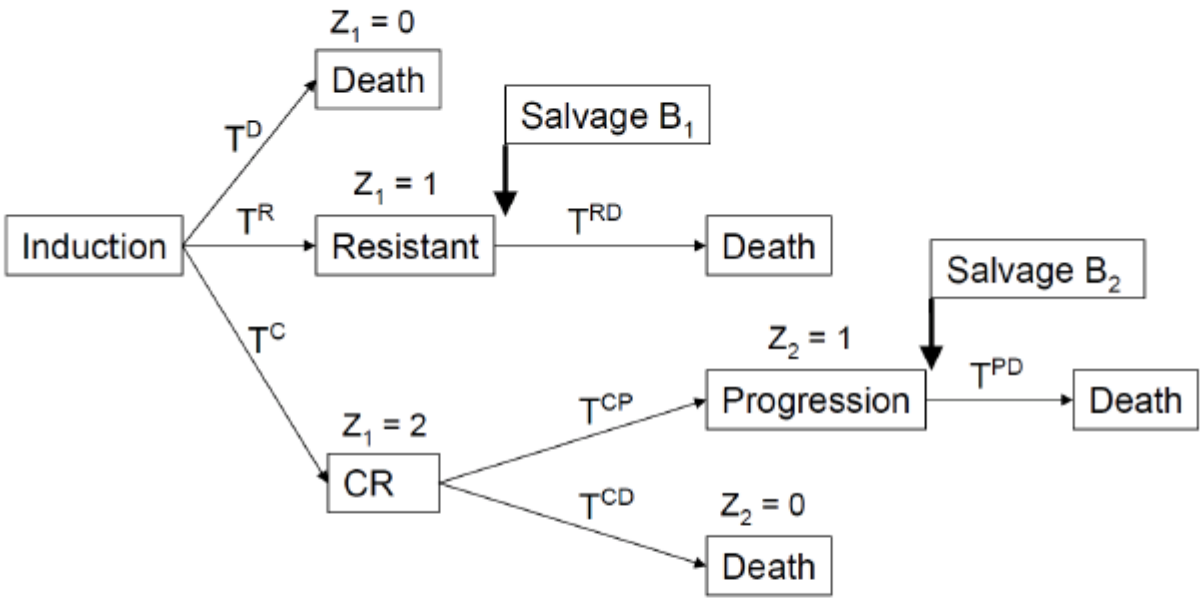

Survival Time =

T^D	if death during induction
$T^R + T^{RD}$	if death after salvage for resistant disease
$T^C + T^{CP} + T^{PD}$	if death after salvage for progression after CR
$T^C + T^{CD}$	if death in CR

The Dynamic Treatment Regime (DTR) denoted by (A, B_1, B_2) says to treat the patient with

- Induction chemo A
- Salvage B_1 if disease is resistant to induction
- Salvage B_2 at disease progression after CR

→ There were 4 x 2 x 2 = 16 possible DTRs

with A from { FAI, FAI+G, FAI + ATRA, FAI+G+ATRA }
and B_1 , B_2 from { HDAC, Other }

Patients were randomized fairly among the induction chemos, but B_1 and B_2 were chosen by the patient's attending physician. → There was selection bias.

Notation for the Bayesian Nonparametric Model

$\boldsymbol{Y^k}$ = log($\boldsymbol{T^k}$) for k^{th} transition time $\boldsymbol{T^k}$ for
k = ***(0, D), (0, C), (0, R), (R, D), (C, D), (C, P), (P, D)***

$\boldsymbol{Z}$ = ($\boldsymbol{Z^1}$, $\boldsymbol{Z^{2,1}}$, $\boldsymbol{Z^{2,2}}$), where $\boldsymbol{Z^1}$ = frontline, $\boldsymbol{Z^{2,1}}$ = salvage if resistant to induction, $\boldsymbol{Z^{2,2}}$ = salvage at progression following CR

$\boldsymbol{x^k}$ = History of all covariates, treatments, transition times up to stage $\boldsymbol{k}$

$[Y^k \mid \boldsymbol{x}^k] \sim F^k(\cdot \mid \boldsymbol{x}^k)$ = probability model for the k^{th} transition time

V_i = observed time of event or censoring for patient i

Product likelihood for all possible observed sequences

$$\mathcal{L} = \prod_{k=1}^{n_T} \prod_{i \in \mathcal{R}^k} f^k(V_i^k \mid \boldsymbol{x}_i^k)^{\delta_i^k} \bar{F}^k(V_i^k \mid \boldsymbol{x}_i^k)^{1-\delta_i^k}$$

$T = \sum_{k=1}^{n_T} T^k$ = Survival time for any of the 4 possible sequences of transition times

Strategy for BNP Model Construction

Construct a BNP survival regression model for each log transition time Y^k = log(T^k) given the entire history

1. Assume a Dependent Dirichlet process (DDP) model for each random distribution G^k by replacing each point mass in G^k with a Gaussian kernel, to extend G^k to a continuous random distribution F^k

2. Endow the Gaussian kernel with regression structure for [$\boldsymbol{Y^k}$ | $\boldsymbol{x^k}$] to account for covariates and outcome history

3. Use the likelihood to define overall survival time as a mixture of the four potential outcomes, T =

$$T^D, \quad T^R + T^{RD}, \quad T^C + T^{CD}, \quad T^C + T^{CP} + T^{PD}$$

The DDP-GP Model for the k^{th} log transition time $Y^k = \log(T^k)$

(as used in in the survival-AUC-CR-Age analysis)

$$F^k(y \mid \boldsymbol{x}^k) = \sum_{h=0}^{\infty} w_h^k \, N(y;\ \theta_h^k(\boldsymbol{x}^k),\ \sigma^k).$$

Gaussian Process (GP) Prior on the Mean of each Normal Summand in the DP Mixture of $F^k(y \mid x^k)$

$$\{\theta_h^k(\boldsymbol{x}^k)\} \sim\ GP(\mu_h^k(\boldsymbol{x}^k), C^k(\boldsymbol{x}^k)), \qquad h = 1, 2, ...$$

$$\mu_h^k(\boldsymbol{x}_i^k) = \boldsymbol{x}_i^k \boldsymbol{\beta}_h^k.$$

$$C^k(\boldsymbol{x}_i^k, \boldsymbol{x}_j^k) = \exp\{-\sum_{m=1}^{M^k}(x_{im}^k - x_{jm}^k)^2\} + \delta_{ij} J^2, \quad i, j = 1, \ldots, n,$$

Empirical Bayes (EB) Prior Specification for each k

For the k^{th} transition time, we assume iid priors $\beta_h^k \sim N(\beta_0^k, \Sigma_0^k)$

$$(\sigma^k)^{-2} \overset{\text{i.i.d.}}{\sim} \text{Ga}(\lambda_1, \lambda_2) \text{ and } \alpha^k \overset{\text{i.i.d.}}{\sim} \text{Ga}(\lambda_3, \lambda_4)$$

To apply the DDP-GP model, one must specify fixed prior hyperparameters $\{\beta_0^k, \Sigma_0^k, k = 1, 2, ...\}$ and $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \lambda_3, \lambda_4)$

1. Do preliminary fit of lognormal for each $Y^k = \log(T^k) \sim N(\boldsymbol{x}^k \beta_0^k, \sigma_0^k)$
2. Use the estimates as the values of the prior means β_0^k
3. Given the empirical estimate of β_0^k tune (λ_1, λ_2) so that the prior mean of σ^k equals it's empirical estimate and its prior var = 1.
4. Set $\lambda_1 = \lambda_2 = 1$ to obtain a vague prior on α^k

Properties of the DDP-GP Model with EB Prior

β_0^k determines the prior mean of the mean function $\{\mu_h^k\}$ of the DP prior, formalizing regression of $\boldsymbol{T^k}$ on $\boldsymbol{x^k}$.

Excessive prior shrinkage might smooth away the treatment effect. The empirical Bayes approach avoids this.

Inference is insensitive to the hyperparameters λ of the priors of the GP variances σ^k and the beta parameters α^k, because :

1) σ^k is the scale of the Gaussian smoothing kernel in the DDP mixture, and has little effect on the imputed fits.

2) α^k determines the number of clusters in the DDP mixture for $F^k(y \mid x^k)$ but most clusters are very small, so it has little effect on the posterior.

The Usual Frequentist IPTW Estimates for Bias Correction

We compute the IPTW estimates for overall mean survival with regime $\mathbf{Z}$ as

$$IPTW(\mathbf{Z}) = \sum_{i=1}^{n} w_i(\mathbf{Z}) T_i \; / \; \sum_{i=1}^{n} w_i(\mathbf{Z}),$$

where

$$\begin{aligned} w_i(\mathbf{Z}) \;=\; & \frac{I(\mathbf{Z}=\mathbf{Z}_i)\delta_i}{\hat{K}(U_i)} \Big[I(s_{1i}=D) + I(s_{1i}=R) I_i(Z^{2,1}) / \hat{\Pr}(Z^{2,1} \mid s_{1i}=R, Z^1, \boldsymbol{x}_i^0, T_i^{(0,R)}) \\ & + I(s_{1i}=C, s_{2i}=D) \\ & + I(s_{1i}=C, s_{2i}=P) I_i(Z^{2,2}) / \hat{\Pr}(Z^{2,2} \mid s_{1i}=C, s_{2i}=P, Z^1, \boldsymbol{x}_i^0, T_i^{(0,C)}, T_i^{(C,P)}) \Big]. \end{aligned}$$

$\hat{K}(U_i)$ = Kaplan-Meier estimate of censoring distribution

The Usual Frequentist AIPTW Estimates for Bias Correction (A = "Augmented")

$$\mathrm{ATE}_{\mathrm{AIPTW}} = \frac{1}{n}\sum_{i=1}^{n}\left\{\left[\frac{I(Z_i=1)Y_i}{\hat{\pi}_i} - \frac{I(Z_i=0)Y_i}{1-\hat{\pi}_i}\right]\right.$$
$$\left. - \frac{I(Z_i=1)-\hat{\pi}_i}{\hat{\pi}_i(1-\hat{\pi}_i)}\left[(1-\hat{\pi}_i)\hat{E}(Y_i \mid Z_i=1,\boldsymbol{x}_i) + \hat{\pi}_i\hat{E}(Y_i \mid Z_i=0,\boldsymbol{x}_i)\right]\right\}$$

where $\hat{\pi}_i$ is the estimated propensity score using logistic regression and $\hat{E}(Y_i \mid Z_i, \boldsymbol{x}_i)$ is estimated by a linear regression model, $i = 0, 1$.

"Double Robustness" of AIPTW estimator :

The AIPTW estimator is consistent if

1) the model for $f(\boldsymbol{Y} \mid \boldsymbol{x})$ is correct or

2) the propensity score model for π is correct.

DDP-GP Estimates of Mean Survival in the Leukemia Trial Data

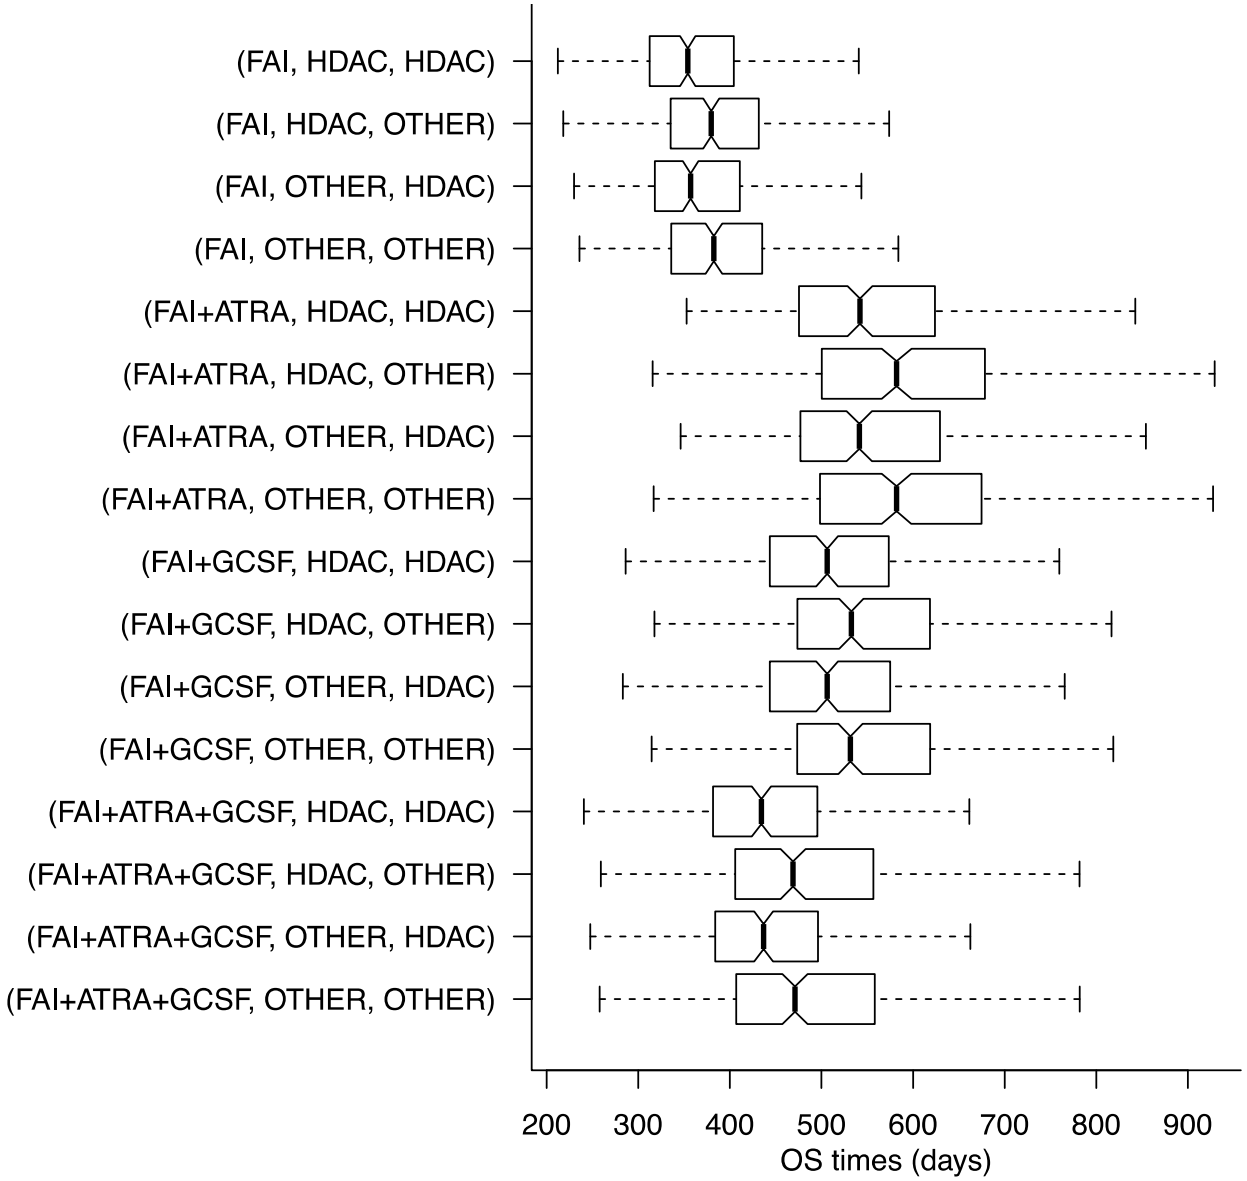

Simulation 1: Two-stage regimes, with covariate-dependent induction and salvage treatment assignments

n=200 patients randomized to induction $Z^1 = a_1$ or a_2 with probabilities depending on the baseline covariate L ~ N(450, 100)

$$p(Z_i^1 = a_1 \mid L_i) = 0.8\, I(L_i < 450) + 0.2\, I(L_i \geq 450).$$

For simplicity, all patients are assumed to be resistant, with transition time $T^{(0,R)}$ ~ lognormal (2 + .0005 L , 0.3)

Survival time $T = T^{(0,R)} + T^{(R,D)}$

Denoting expit(u) = $e^u/(1+e^u)$, salvage treatments Z^2 assigned with probabilities

$$p(Z_i^2 = 1 \mid Z_i^1, T_i^{(0,R)}) = Z_i^1 \text{expit}(1 - 0.003\, T_i^{(0,R)}) + (1 - Z_i^1)\text{expit}(-0.8 - 0.004\, T_i^{(0,R)})$$

$$T_i^{(R,D)} \sim \text{LN}(\boldsymbol{\beta}^{(R,D)} \boldsymbol{x}_i^{(R,D)}, 0.3), \text{ where } \boldsymbol{\beta}^{(R,D)} = (-0.5,\ 0.03,\ 0.2, 0.5,\ 0.3)$$

$$\boldsymbol{x}_i^{(R,D)} = \quad (1,\ L_i,\ Z_i^1, \log(T_i^{(0,R)}), Z_i^2)$$

Goal: Estimate mean survival for each DTR = (Z^1, Z^2)

Results of Simulation 1 : AIPTW and IPTW use the correct (simulation truth) model for the outcomes

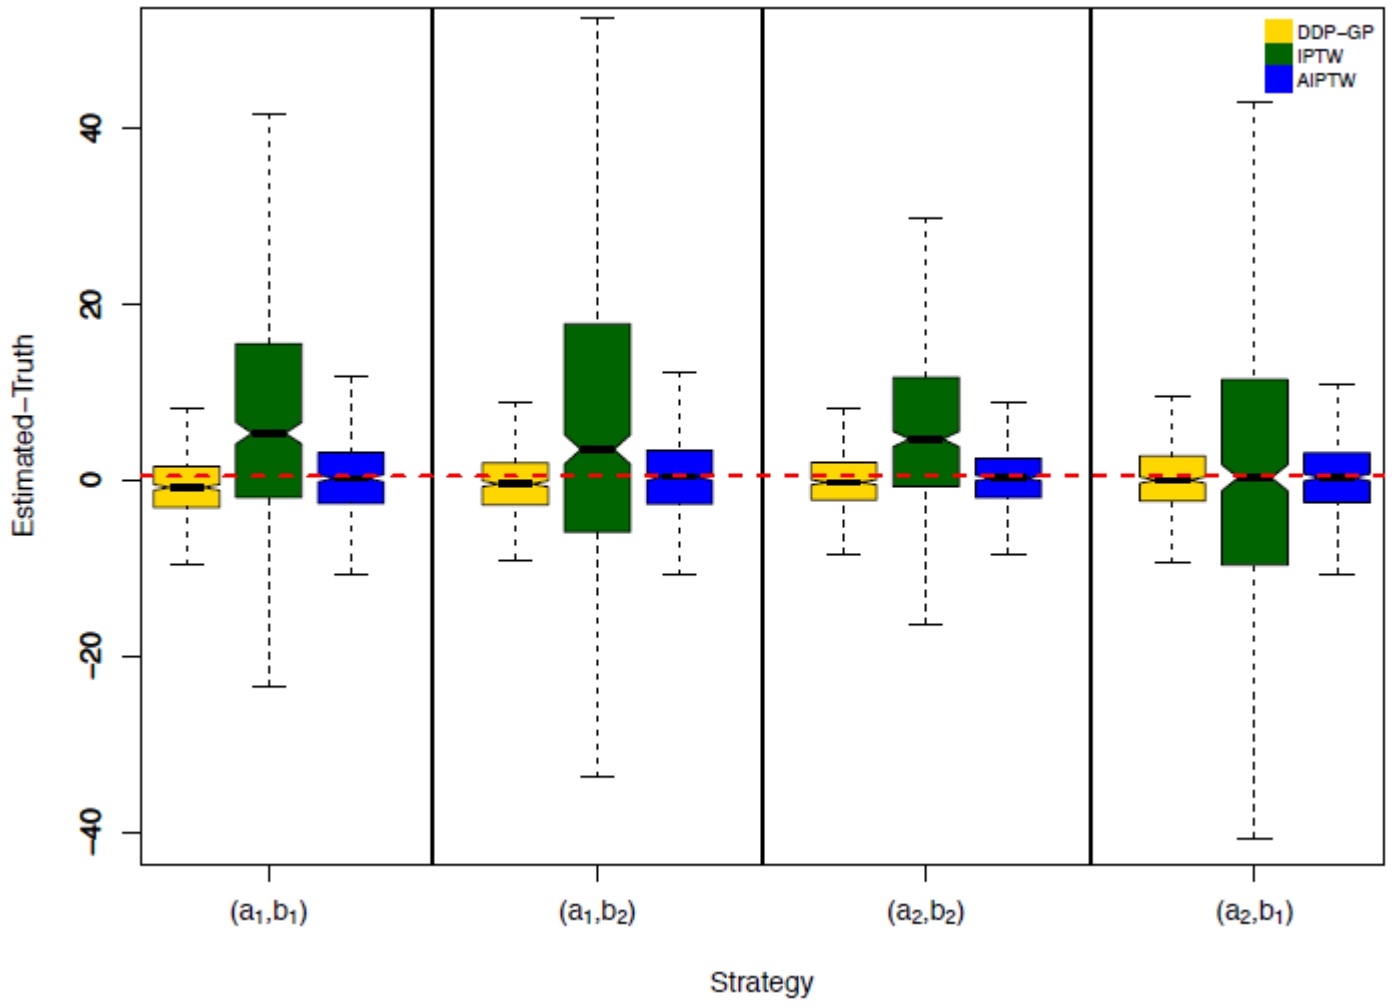

Simulation 2: A stylized version of the leukemia data
AIPTW does not use the correct (simulation truth) model

Multi-stage regimes with covariate-dependent induction and salvage

n=200 patients randomized fairly to induction $Z^1 = a_1$ or a_2 with probabilities .50 each. 15% censoring.

Baseline covariate L ~ N(100, 100)

- Resistant patients randomized between salvage b_{11} or b_{12} with probability

$$p(Z^{2,1} = b_{11} \mid L_i) = 0.8\, I(L_i < 100) + 0.2\, I(L_i \geq 100).$$

- Patients who achieved CR and then progressed, (C, P), were randomized between salvage b_{21} or b_{22} with probability

$$p(Z^{2,2} = b_{21} \mid L_i) = 0.2\, I(L_i < 100) + 0.85\, I(L_i \geq 100).$$

Simulation 2: A stylized version of the leukemia data

$T^{(0,R)}$ and $T^{(0,C)}$ simulated as lognormals depending on $\boldsymbol{x} = (L , Z^1)$

For k = (R,D), (C,P), (P,D), the transition times were simulated as lognormals depending on each entire history

$$\boldsymbol{x}^{(R,D)} = (L , Z^1 , T^{(0,R)} , Z^{2,1})$$

$$\boldsymbol{x}^{(C,P)} = (L , Z^1 , T^{(0,C)})$$

$$\boldsymbol{x}^{(P,D)} = (L , Z^1 , T^{(0,C)} , T^{(C,P)} , Z^{2,2})$$

$$T_i = \begin{cases} T_i^{(0,R)} + T_i^{(R,D)} & \text{if patient } i \text{ had sequence } (L, Z^1, T^{(0,R)}, Z^{2,1}) \\ T_i^{(0,C)} + T_i^{(C,P)} + T_i^{(P,D)} & \text{if patient } i \text{ had sequence } (L, Z^1, T^{(0,C)}, T^{(C,P)}, Z^{2,2}). \end{cases}$$

Goal: Estimate mean survival for each DTR = (Z^1, $Z^{2,1}$, $Z^{2,2}$)

Results of Simulation 2

AIPTW and IPTW do not use the correct (simulation truth) model

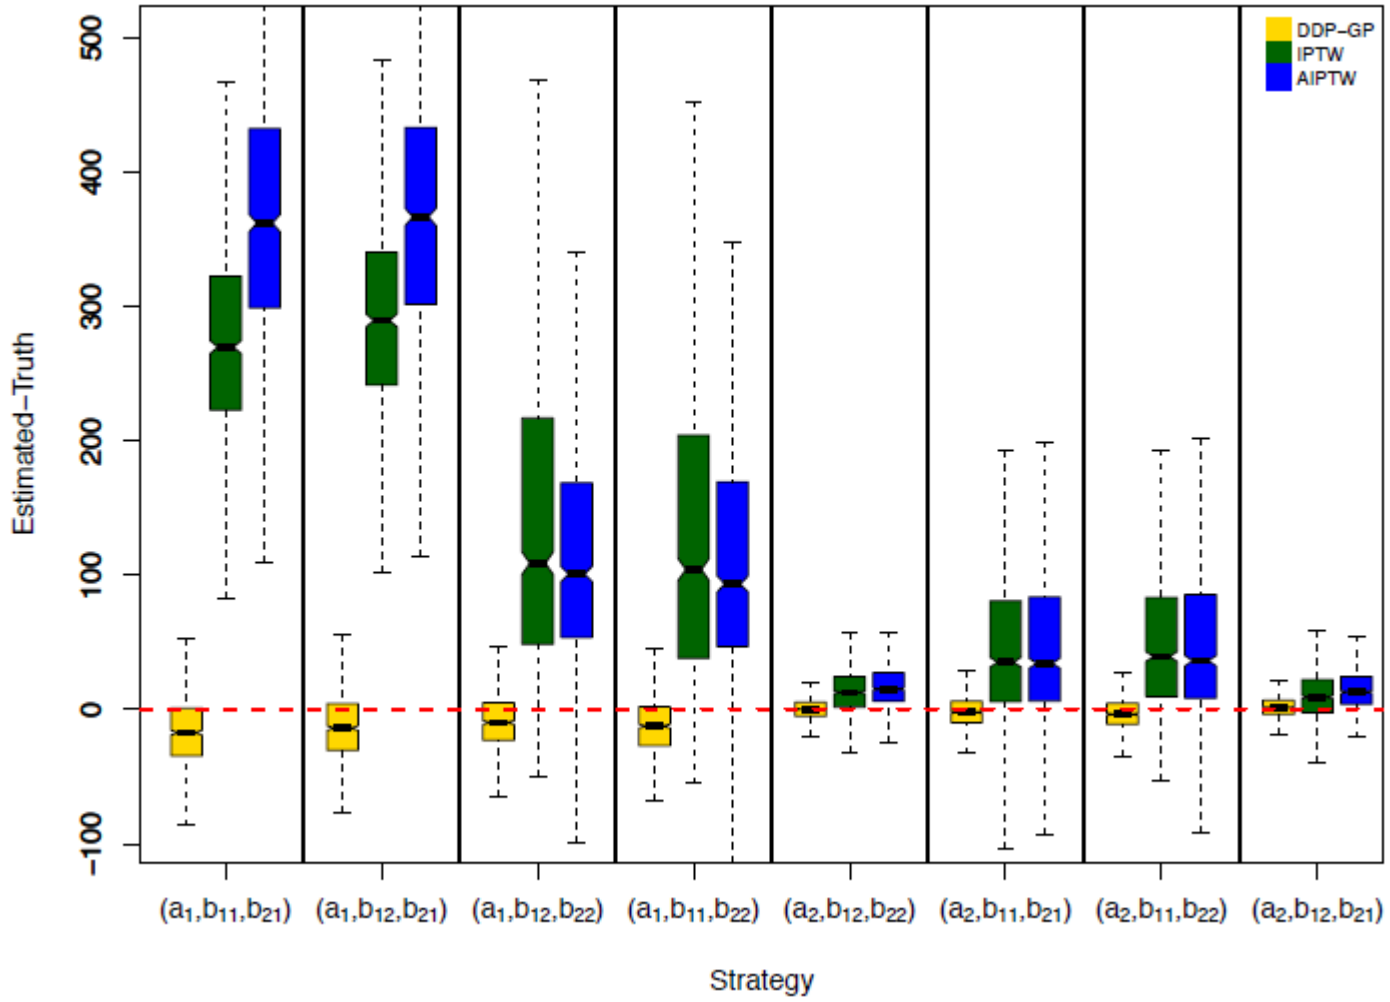

Additional References not in the *Pharma Stat* paper

Mueller P and Quintana FA. (2004). Nonparametric Bayesian data analysis. *Statistical Science* 19:95-110, 2004.

Mueller P and Mitra R. Bayesian nonparametric inference - Why and how. *Bayesian Analysis* 8:269-232, 2013.

Xu Y, Thall PF, Mueller P and Mehran, RJ. A decision-theoretic comparison of treatments to resolve air leaks after lung surgery based on nonparametric modeling. *Bayesian Analysis* 12:639-652, 2017.

Xu Y, Thall PF, Hua W and Andersson B. Bayesian nonparametric survival regression for optimizing precision dosing of intravenous busulfan in allogeneic stem cell transplantation. *J Royal Statistical Society, Series C* 68:809-828, 2019.