Brian Tom

Personalising inter-donation intervals amongst blood donors

12th November 2019

Centre for Statistical Methodology Symposium, London

Joint work with Yuejia Xu, Angela Wood and Michael Sweeting

What is this talk about?

- Describing how we approach personalisation of blood donation strategies (as a concrete example)
- Highlighting some of the generic issues/challenges to confront when adopting a personalise medicine approach

Biology and Medicine has moved on

- New Era
  - Explosion of data sources and types now available for understanding biological processes, diseases and informing health, diagnosis and care

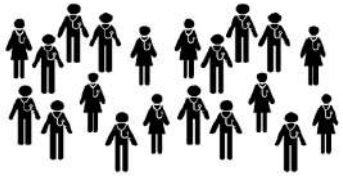

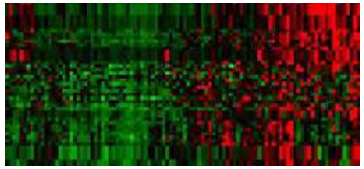

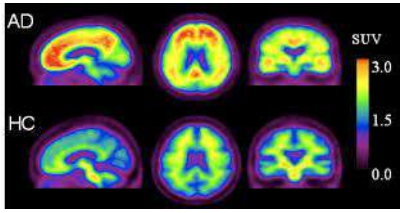

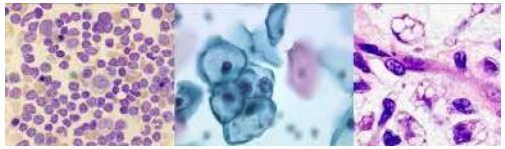

Personalised Medicine (PM)

- Not surprising to use this extensive information and data from patient cohorts, healthy populations, RCTs and experimental systems to better understand and characterise disease

- Then use findings to target treatment and make more informed decisions according to characteristics (biological or otherwise) of the patient shared with a subgroup of "*similar*" individuals

Stratified or Precision Medicine (PM)

- What's new is the…
  - Resolution or granularity (towards ever finer stratification)
  - Complexity and information content of the data used to stratify (data coming "*closer*" to matching disease complexity)
  - Scale ("big data")
  - Stratification to facilitate understanding (causal questions, guide therapeutic development)

PM perspectives (wrt treatment)

- High patient heterogeneity → Significant heterogeneity in treatment response
  - "One-size-fits-all" approach is suboptimal

- Typically conceptualised under two perspectives (frameworks)

1. Right patient (or patient subgroup) for a given treatment

   Subgroup detection problem (reflect pharma's view)

2. Right treatment (from various options) for a given patient

   Estimation of optimal individualise treatment rules (ITRs)
   (reflect more society's/patient's or policy maker's view)

Lipkovich et al. (2017) Tutorial in biostatistics: data-driven subgroup identification and analysis in clinical trials. Statistics in Medicine, 36(1): 136-196

Motivation – INTERVAL Trial

- In UK, men donate blood a minimum of every 12 weeks; women every 16 weeks
- Limits exists to safeguard donor health
- Never been a large scale study whether or not these standard inter-donation intervals are optimal for ensuring *donor health* while maintaining adequate *blood supplies*
- INTERVAL – 1st randomised trial to determine whether blood can be *safely* and *acceptably collected* from donors by NHSBT more frequently than present practice
  (Moore et al. (2014); Trials 15(1):363)

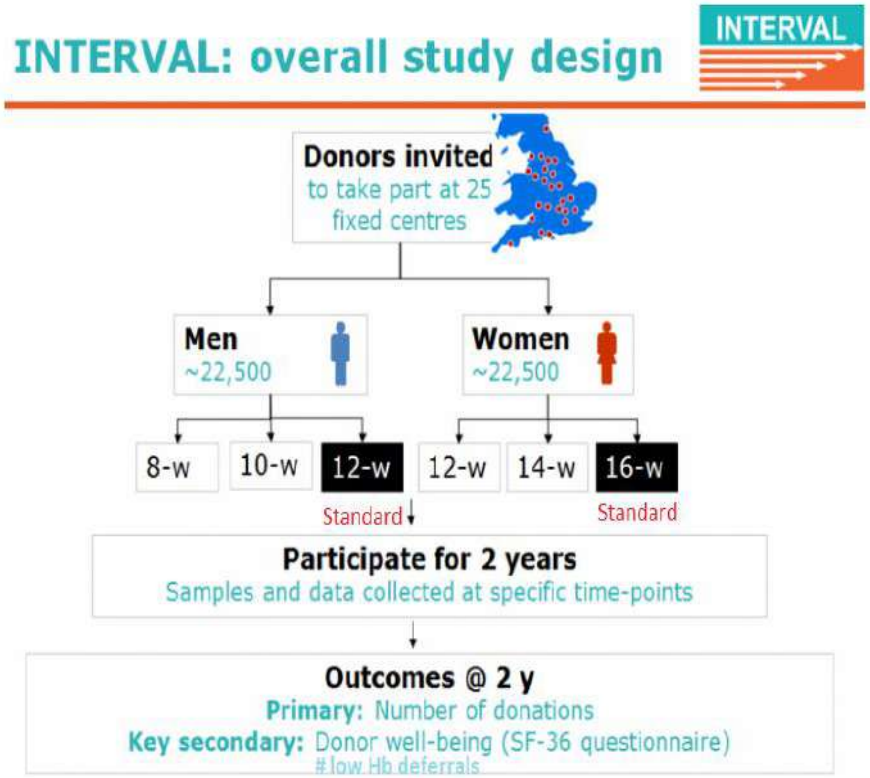

- Main findings of INTERVAL:
  More frequent donation resulted in *substantially more* blood collected, *without* major impact on QoL, physical activity, or cognitive function. However, *more* donation-related symptoms, deferrals and iron deficiency. (Di Angelantonio et al. (2017); Lancet 390, 2360-71)

PM approach to the INTERVAL Trial

- Can we develop evidence-based tools (strategies) that identify people capable of "high intensity" donations and more vulnerable to donation-related adverse events?

- Clearly, a PM question!!!
  - Define *individualised treatment rule* (ITR), D(**X**), that maps a donor's covariate vector, **X**, to one of the set of available treatment options, *A*
  - Choose the rule that optimizes the expected outcome, *Y*

- Tailor *donation frequency* to donor's *donation capacity* based on donor-specific *baseline* characteristics
  - Estimate optimal individualised inter-donation intervals
  - Outcomes: Total whole units of blood over 2 years (1 unit = 470ml), # low Hb deferrals (min of 3 mths before allowed to donate again)
  - Baseline covariates: age, BMI, ethnicity, blood group, SF36 PCS, MCS, donor status, units of blood donation in prior 2-years, smoking, alcohol consumption
  - Routine blood measurements: Hb level, WBC count, RBC count, mean corpuscular haemoglobin, mean corpuscular volume and platelet count
  - Blood-based biomarkers: ferritin and transferrin

What's nice about using INTERVAL for PM

- Size: approx. 22,500 males and 22,500 females
  - Can consider using both statistical and machine learning approaches

- Randomised controlled trial and a single decision point (baseline)
  - Assumptions required to construct the counterfactuals more readily satisfied
  - Consistency: observed outcome same as counterfactual/potential outcome under treatment actually received
  - No Unmeasured Confounders: treatment assignment indep. of potential outcomes conditional on baseline covariates
  - Positivity assumption: wp 1, propensity score is bounded away from 0. All subjects had a chance of receiving any particular treatment. This is true for INTERVAL after stratifying by gender

- Many interesting features that add more complexity
  - Multiple treatments (3) and ordering
  - Multiple outcomes: benefit and risk

How should one proceed?

- Number of methods proposed in literature to estimate the optimal ITR
  - Divided into indirect and direct methods
  - Typically developed for only two treatments

- Indirect methods
  - Postulate a model for the conditional mean of the response given the covariates and treatment, E[Y | $\mathbf{X}$=$\mathbf{x}$, A=a], say a regression model of the form h($\mathbf{x}$) + c($\mathbf{x}$)a; and
  - Decision inferred by maximising the outcome from the fitted model (i.e. arg max$_a$ $\widehat{\mathrm{E}}$[Y | $\mathbf{X}$=$\mathbf{x}$, A=a]); or
  - Modelling contrast in conditional means between two treatment options and
  - Decision between the two treatment options is based on sign of the estimated contrast

- Direct methods
  - Directly consider the value function, V(D($\mathbf{X}$)), which is the expected potential outcome for a specific decision rule (i.e. E[Y^*(D($\mathbf{X}$)] = E[$\frac{I(A=D(\mathbf{X}))}{P(A|\mathbf{X})}Y$])
  - Optimal ITR is the one that maximizes the expected potential outcome among the class of all possible regimes if applied to the population of subjects under study
  - Becomes a weighted classification problem; machine learning can be considered

How should one proceed?

- For the INTERVAL Trial, considered methods that are scalable, do not impose high computational costs, use complexity control (favour parsimony) and ignore treatment ordering

- Indirect methods considered were
  - L1-penalized least squares; regression approach (Qian and Murphy (2011), using group or hierarchical group lasso)
  - Bayesian additive regression trees; non-parametric, ensemble approach (Chipman et al., 2010)
  - Note that treatment-by-covariate handled differently in these two approaches

- Direct methods considered were
  - Adaptive contrast weighted learning; weighted classification problem based on weights corresponding to the lower or upper bound on the expected loss in the outcome due to suboptimal treatments (Tao and Wang, 2017)
  - Multi-category D-learning; uses regressions for either linear or non-linear decision rules (such as WLS with or without lasso) to solve multiple pairwise comparison problems. Each treatment effect is then based on averaging estimated pairwise contrasts involving that treatment. Treatment with largest effect is chosen as optimal (Qi and Liu, 2018)

Standard approach

- Look for treatment-by-covariate interactions!
- Di Angelantonio et al. (2017) had a number of subgroup analyses that were pre-specified
- Subgroups findings indicate the presence of quantitative interactions
- Conclude: Evidence for personalising inter-donation intervals (?)
- Not correct! Quantitative interactions are unhelpful for determining optimal ITRs!
- Direction of treatment effect is consistent across all values of covariate

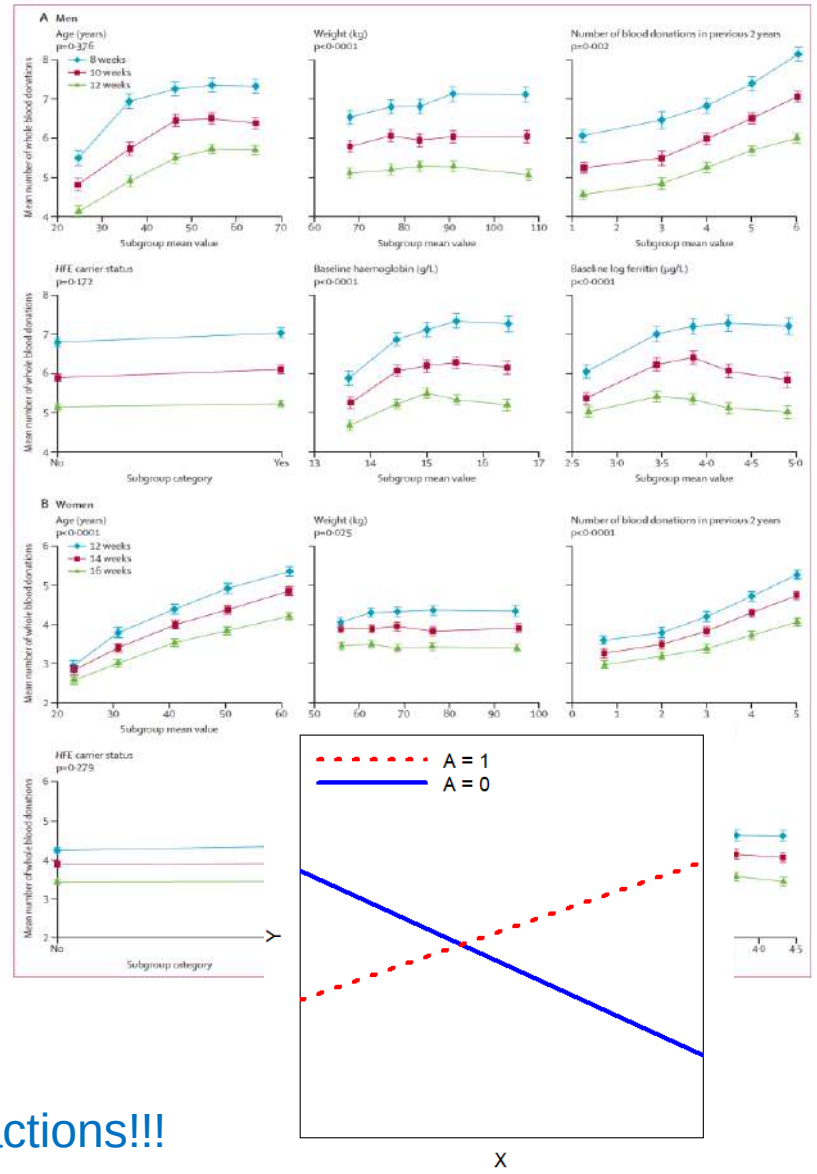

Qualitative Interactions!!!

Results for males (based on 100 rep of 5-fold CV)

Target Outcome	Method	Assignment Percentages			ITR Effects	
		12 weeks	10 weeks	8 weeks	Donation	Deferral
Donation	l_1-PLS-HGL	0.1 (0.0)	0.3 (0.0)	99.6 (0.0)	1.308 (0.004)	0.026 (0.000)
	l_1-PLS-GL	0.0 (0.0)	0.3 (0.3)	99.7 (0.3)	1.311 (0.005)	0.027 (0.000)
	ACWL	0.0 (0.0)	0.0 (0.0)	100.0 (0.0)	1.315 (0.000)	0.027 (0.000)
	D-learning	0.3 (0.1)	0.3 (0.2)	99.4 (0.2)	1.307 (0.006)	0.027 (0.000)
Deferral	l_1-PLS-HGL	94.4 (0.6)	5.5 (0.6)	0.0 (0.0)	-1.188 (0.010)	-0.024 (0.000)
	l_1-PLS-GL	99.7 (0.6)	0.2 (0.5)	0.0 (0.1)	-1.246 (0.006)	-0.024 (0.000)
	ACWL	99.7 (0.6)	0.3 (0.6)	0.0 (0.0)	-1.244 (0.007)	-0.025 (0.000)
	D-learning	95.7 (0.5)	4.1 (0.5)	0.2 (0.1)	-1.200 (0.011)	-0.024 (0.000)

- ITR effects correspond to either average increase in blood donation per donor over 2 years; or increased number of low Hb deferrals per 100 donor-attendances, on average

- Strong indication that in the extreme cases where the true optimal decision rules are trivial
  - Non-personalised/fixed
  - Quantitative interactions only present? (Some evidence based on L1-PLS)

- Maximize total units of blood – assign donors to 8 weeks;
  Minimize low Hb deferrals – assign donors to 12 weeks (polar opposite)

- BART: 100% to 8wks,1.313, 95% CrI [1.237-1.392]; 91.6% to 12 wks, -0.024 [-0.027,-0.022]

Utility (composite) outcome

- Do findings change if we consider an outcome which balances benefit and risk (i.e. a utility)?

- Define the utility as U = Total units of blood – b x No. of low Hb deferrals
  - Discounted units of blood collected by increased occurrence of Hb deferrals
  - b is trade-off parameter reflecting equivalent benefit loss (in total units of blood) for 1 unit increase in risk (one extra deferral) per donor over 2 years
  - b ideally should be elicited from subject matter experts (e.g. b=2, NHSBT)

		Assignment Percentages			ITR Effects		
Trade-off Parameter	Method	12 weeks	10 weeks	8 weeks	Donation	Deferral	Utility
$b = 1$	l_1-PLS-HGL	0.9 (0.1)	1.2 (0.1)	97.9 (0.1)	1.309 (0.008)	0.024 (0.001)	1.064 (0.009)
	l_1-PLS-GL	0.3 (0.1)	2.4 (0.8)	97.2 (1.0)	1.289 (0.014)	0.025 (0.001)	1.040 (0.014)
	ACWL	0.0 (0.0)	0.0 (0.1)	100.0 (0.1)	1.314 (0.003)	0.027 (0.000)	1.055 (0.002)
	D-learning	0.7 (0.2)	1.2 (0.5)	98.1 (0.5)	1.309 (0.012)	0.025 (0.001)	1.058 (0.013)
$b = 2$	l_1-PLS-HGL	3.4 (0.1)	4.2 (0.3)	92.4 (0.4)	1.242 (0.016)	0.021 (0.001)	0.809 (0.020)
	l_1-PLS-GL	1.7 (0.4)	7.2 (1.6)	91.1 (2.0)	1.217 (0.027)	0.022 (0.002)	0.774 (0.024)
	ACWL	2.7 (0.8)	3.4 (1.7)	93.9 (1.3)	1.266 (0.019)	0.022 (0.001)	0.814 (0.016)
	D-learning	1.5 (0.4)	5.8 (1.1)	92.7 (1.0)	1.260 (0.022)	0.022 (0.001)	0.816 (0.023)
$b = 3$	l_1-PLS-HGL	8.6 (0.3)	11.9 (0.5)	79.5 (0.6)	1.091 (0.022)	0.011 (0.001)	0.689 (0.028)
	l_1-PLS-GL	4.8 (1.1)	15.0 (3.3)	80.2 (4.4)	1.069 (0.056)	0.017 (0.003)	0.569 (0.041)
	ACWL	9.3 (1.4)	8.2 (2.7)	82.5 (2.2)	1.100 (0.034)	0.014 (0.001)	0.627 (0.032)
	D-learning	3.8 (0.8)	17.5 (1.6)	78.6 (1.3)	1.067 (0.027)	0.016 (0.001)	0.607 (0.027)
$b = 4$	l_1-PLS-HGL	17.0 (0.4)	23.3 (0.5)	59.7 (0.5)	0.745 (0.023)	0.001 (0.001)	0.623 (0.030)
	l_1-PLS-GL	10.5 (2.3)	27.9 (3.2)	61.6 (5.2)	0.782 (0.070)	0.008 (0.004)	0.468 (0.081)
	ACWL	16.8 (1.9)	16.2 (3.5)	67.0 (3.1)	0.793 (0.055)	0.007 (0.002)	0.475 (0.033)
	D-learning	9.9 (1.6)	30.0 (1.7)	60.2 (0.7)	0.783 (0.027)	0.006 (0.001)	0.543 (0.039)
$b = 5$	l_1-PLS-HGL	26.4 (0.4)	33.4 (0.5)	40.3 (0.3)	0.410 (0.022)	-0.007 (0.001)	0.648 (0.031)
	l_1-PLS-GL	18.2 (3.2)	48.4 (6.2)	33.4 (3.4)	0.324 (0.059)	-0.004 (0.002)	0.485 (0.084)
	ACWL	30.1 (2.5)	22.6 (4.2)	47.3 (3.2)	0.422 (0.053)	-0.005 (0.002)	0.541 (0.046)
	D-learning	19.3 (1.6)	37.4 (1.3)	43.3 (0.7)	0.505 (0.031)	-0.004 (0.001)	0.622 (0.045)

Utility (composite) outcome

- Do findings change if we consider an outcome which balances benefit and risk (i.e. a utility)?
  - Yes, becomes less extreme and indicates substantive treatment effect heterogeneity

- What about non-personalized/fixed rules?

	ITR Effects on Utility				
Non-personalized Rule	$b = 1$	$b = 2$	$b = 3$	$b = 4$	$b = 5$
Recommend all male donors to donate every 12 weeks	-1.308	-0.828	-0.618	-0.408	-0.199
Recommend all male donors to donate every 10 weeks	-0.025	0.027	0.079	0.131	0.183
Recommend all male donors to donate every 8 weeks	1.055	0.795	0.535	0.275	0.015

- Conclude: added value in personalizing over adopting one-size-fits-all approach for the utility

- Note: despite probable difference in form of estimated rules, optimal ITRs estimated across different approaches for a given donor can still overlap (b=1, 2 or 3; substantial agreement)

- Utility outcome can be extended to incorporate other aspects important to decision-making (e.g. use of novel biomarkers, retention of specifically important donor sub-populations, costs)

- Alternatively, handle trade-off between benefit and risk from constrained optimization framework
  - maximize benefit under constraint of controlling average risk under a meaningful threshold

Ordinal treatment

- By ignoring ordinality, expect to lose information potentially helpful in improving prediction performance of estimated ITRs. Results in suboptimal decision making
  - Ordering of treatments may indicate that there is also an ordering of the loss incurred
  - A suboptimal treatment closer to true optimal would be expected to have a smaller loss in outcome than those further from the truth
  - For example, for a female donor whose true optimal inter-donation interval is 16 weeks, incorrectly allocating her to 12-week inter-donation interval may lead to more severe consequences on donor health than to the 14-week option
  - Similarity with dose-finding problems
- We propose a method to estimate optimal ITRs here, which takes advantage of the ordering
- Basic idea:
  - Decompose ordinal problem into multiple binary comparison subproblems (classification)
  - Aggregate multiple binary decisions in order to construct decision tree
- We call the method sequential re-estimation (SR) learning
  - Includes a series of sequential steps determining whether a more intensive treatment should be started
  - Followed by re-estimation steps that compare two consecutive treatments to reconsider for a step-up in treatment if in sequential step, the decision made was to not intensify treatment

Sequential re-estimation (SR) learning

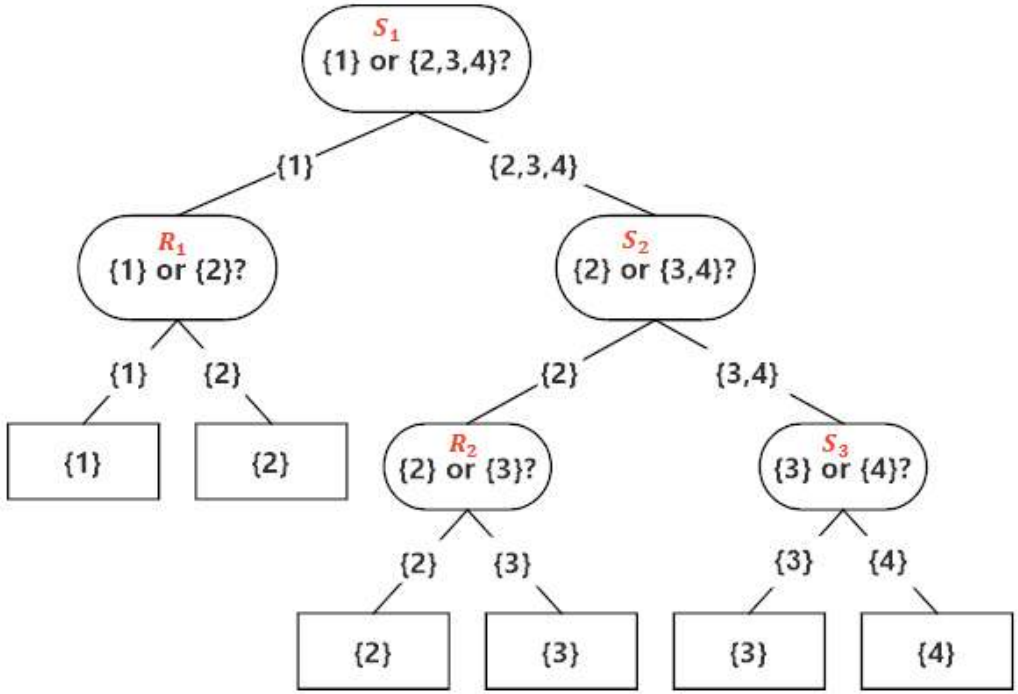

- Method handles linear and non-linear decision boundaries
- Binary subproblems solved by weighted SVM with or without variable selection; uses convex optimization
- Simulations demonstrate superior performance over methods not taking account of ordering

Application to an “of-interest” INTERVAL subpop’n

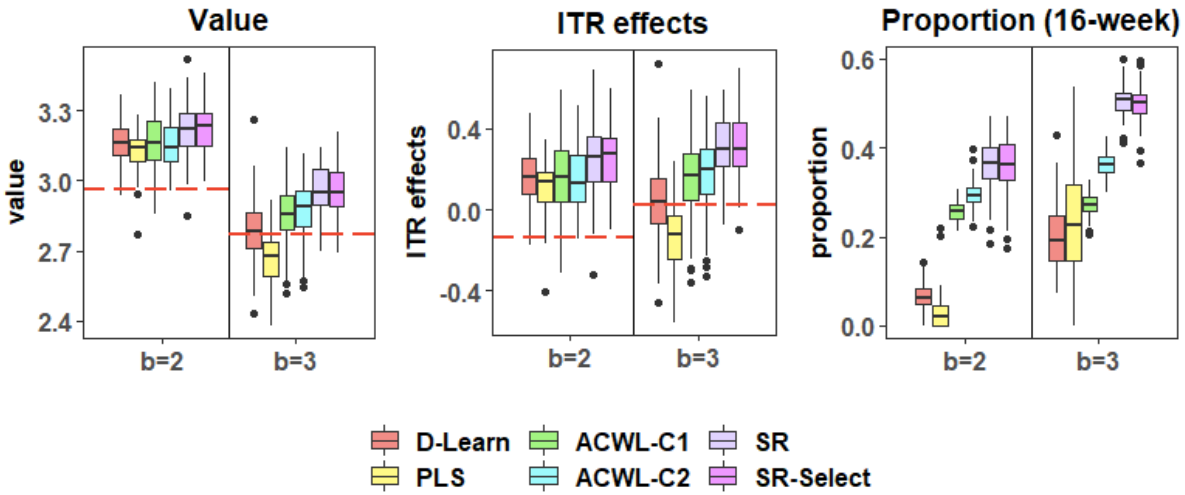

- Target donor subpopulation of n=884 females, aged < 40 with O- blood type (valuable)
- Target outcome is the utility, U = units of blood collected – b x no. of low Hb deferrals (b=2 or 3)
- Performance evaluated using 100 repetitions of 5-fold CV
  - Empirical value function
  - ITR effects: mean difference in outcome between optimal ITR followers and non-followers
  - % assigned to 16-week inter-donation interval (safest one and current practice)
- Proposed SR learning (with and without variable selection) estimated to have the largest value, ITR effects and proportion allocated to 16-weeks (i.e. less aggressive)

Summary

- Presented how one can go about learning optimal decision rules in the setting of a large RCT with mapping of baseline covariates to one of multiple treatments
- Demonstrated that even in the setting where “one-size-fits-all” approach is the truth, methods developed for estimating optimal treatment rules seem to converge to the correct non-personalized/fixed rule
- Highlighted the fact that what is required are qualitative interactions not quantitative
- Discussed how to handle multiple outcomes and ordinality of treatments
- Did not discuss the other PM perspective (i.e. right patient/subgroup for a given treatment), although relevant here (e.g. identify a subpopulation of “super donors” who have a capacity for attending every 8 weeks)
- Nor discussed extension to dynamic treatment regimes or determining when to return based on prediction of the recovery of Hb levels to safe levels (see Nasserinejad et al. 2015)
- Finally although not discussed, validation is a crucial aspect of PM

References

- Lipkovich, Ilya, Alex Dmitrienko, and Ralph B D'Agostino Sr (2017). "Tutorial in biostatistics: data-driven subgroup identification and analysis in clinical trials." Statistics in medicine 36(1): 136-196.
- Moore, Carmel, et al. (2014). "The INTERVAL trial to determine whether intervals between blood donations can be safely and acceptably decreased to optimise blood supply: study protocol for a randomised controlled trial." Trials, 15(1): 363
- Di Angelantonio, Emanuele, et al. (2017). "Efficiency and safety of varying the frequency of whole blood donation (INTERVAL): a randomised trial of 45 000 donors." The Lancet, 390(10110): 2360-2371.
- Qian, Min, and Susan A. Murphy (2011). "Performance guarantees for individualized treatment rules." Annals of Statistics, 39(2): 1180.
- Chipman, Hugh A., Edward I. George, and Robert E. McCulloch (2010). "BART: Bayesian additive regression trees." The Annals of Applied Statistics 4(1): 266-298.
- Tao, Yebin, and Lu Wang (2017). "Adaptive contrast weighted learning for multi-stage multi-treatment decision-making." Biometrics 73(1): 145-155.
- Qi, Zhengling, and Yufeng Liu (2018). "D-learning to estimate optimal individual treatment rules." Electronic Journal of Statistics 12(2): 3601-3638
- Nasserinejad, Kazem, et al. (2013). "Predicting hemoglobin levels in whole blood donors using transition models and mixed effects models." BMC Medical Research Methodology 13(1): 62.