A Forward-looking Covariate-adjusted Response-adaptive Patient Randomisation Rule:

Delivering Personalised Treatments, Patient Benefit and Power Gains in Multi-armed Trials

Sofía S. Villar

MRC | Biostatistics Unit

UNIVERSITY OF CAMBRIDGE

Cambridge, 21st June 2018

BAYES 2018: Bayesian Biostatistics Workshop

This talk is based on joint work with Prof. W. Rosenberger (GMU)

Outline

Tailoring Treatment to a Patient's Characteristics

Testing Several Arms in Few Patients

- FDA approved several cancer drugs for use in patients whose tumours have specific genetic characteristics (identified by a diagnostic test). Also, for cystic fibrosis patients with a specific genetic mutation.
- This has strengthened the promise of "*personalised medicine*" - the tailoring of treatment to the individual characteristics of each patient.
- The challenge is how can we answer which patients respond differently to a treatment in contexts in which there are several promising new treatments and relatively few patients to test them - even fewer patients if a treatment works only within a subgroup.
- Research for personalised medicine requires methodology for trial design to identify superior treatments among several arms more quickly, mainly treatments that work better within subgroups.

Response-adaptive Randomisation and Covariates

Assigning More Patients to Better Treatment based on their Profile

- Response-adaptive randomisation (RAR) alters the allocation probabilities based on observed outcomes up to that patient.

  The usual goal is to achieve some *ethical* or *statistical objective*.

- Covariate-adjusted Response-adaptive (CARA) alters the allocation probabilities given observed outcomes and covariates and the current patient's covariate profile.

  CARA designs can benefit patients (in case of a treatment-covariate interaction) by skewing the allocation probabilities to the better performing treatment among patients with a similar covariate profile.

! Note that CARA early covariate-adaptive methods (1970s) try to induce balance between treatment groups with respect to covariate margins: Pocock and Simon (1975); Taves, D. R. (1974); Wei, L.J. (1978). Balance does not lead to efficiency or ethically attractive designs (Rosenberger and Sverdlov, 2008).

Limitations of Existing CARA Procedures

Developing Efficient Non-myopic Randomised CARA Procedures

(1) Almost all RAR procedures are myopic - be they Bayesian or not - adapt based on past data only and not potential future data.

Forward looking RAR procedures are computationally expensive and deterministic. Recent work proposed a randomised tractable non-myopic RAR algorithm using *bandit* theory. Villar et al (2015)

(2) RAR procedures with ethical goals have considerably higher variability (low power) and other estimation issues (Villar et al , 2015).

E.g. CARA procedure Rosenberger, Vidyashankar and Agarwal (2001)

- In Villar and Rosenberger (2018) we develop a non-myopic CARA procedure with patient benefit advantages which is:

computationally feasible (implementable);
randomised (selection bias control);
efficient in the multi-armed case (high power);

Outline

The Multi-armed Bandit Problem and the Gittins Index

Introducing Randomization to Bandit Strategies

- The multi-armed bandit problem (MABP): optimal allocation of treatments to patients with an overall patient benefit goal.
- The Gittins Index (GI) (Gittins and Jones, 1979) is a tractable rule that recovers the (intractable) optimal solution to the MABP.

  Both the MABP and the GI use a Bayesian updating procedure.
- Despite its tractability and considerable ethical advantages a main barrier to its use is the lack of randomisation of treatment allocations.
- Villar et al (2015) defined a block randomised procedure based on the GI. FLGI allocation probabilities: $\pi_{k,j}$, the probability of allocation to treatment k at stage j - common to all patients in block j- when using the GI and given observed responses up to block j-1.

The Forward Looking Gittins Index

Introducing Randomization to the Gittins Index Rule

Assume that T patients arrive sequentially in blocks of size b over J stages, so that $J \times b = T$. In Villar et al (2015) we defined group allocation probabilities based on the Gittins Index (GI) rule as follows:

Simplest example: $b = 2$. Priors: control $(s_{(0,0)}, f_{(0,0)}) = (1, 2)$ and experimental $(s_{(1,0)}, f_{(1,0)}) = (1, 1)$

$$j = 1,$$
$$\mathcal{G}_1(1,1) = 0.8699$$
$$\mathcal{G}_0(1,2) = 0.7005$$

$\frac{1}{2}$ / $Y_{1,0} = 1$

$$j = 2,$$
$$\mathcal{G}_1(2,1) = 0.9102$$
$$\mathcal{G}_0(1,2) = 0.7005$$

$\frac{1}{2}$ / $Y_{1,0} = 0$

$$j = 2,$$
$$\mathcal{G}_1(1,2) = 0.7005$$
$$\mathcal{G}_0(1,2) = 0.7005$$

What is the (patient-average) probability of each arm being allocated in the next block using the GI (and given the priors)?

$$\pi_{1,0} = \frac{(0 \times 1) + (0 \times 1/2 + 1/2 \times 1/2)}{2} = 1/8 \ , \ \pi_{1,1} = \frac{(1 \times 1) + (1 \times 1/2 + 1/2 \times 1/2)}{2} = 7/8.$$

The Forward Looking Gittins Index

Introducing Randomization to the Gittins Index Rule

Assume that T patients arrive sequentially in blocks of size b over J stages, so that $J \times b = T$. In Villar et al (2015) we defined group allocation probabilities based on the Gittins Index (GI) rule as follows:

Simplest example: $b = 2$. Priors: control $(s_{(0,0)}, f_{(0,0)}) = (1, 2)$ and experimental $(s_{(1,0)}, f_{(1,0)}) = (1, 1)$

$$j = 1,$$
$$\mathcal{G}_1(1,1) = 0.8699$$
$$\mathcal{G}_0(1,2) = 0.7005$$

$\frac{1}{2}$ / $Y_{1,0} = 1$

$$j = 2,$$
$$\mathcal{G}_1(2,1) = 0.9102$$
$$\mathcal{G}_0(1,2) = 0.7005$$

$\frac{1}{2}$ / $Y_{1,0} = 0$

$$j = 2,$$
$$\mathcal{G}_1(1,2) = 0.7005$$
$$\mathcal{G}_0(1,2) = 0.7005$$

What is the (patient-average) probability of each arm being allocated in the next block using the GI (and given the priors)?

$$\pi_{1,0} = \frac{(0 \times 1) + (0 \times 1/2 + 1/2 \times 1/2)}{2} = 1/8 \ , \ \pi_{1,1} = \frac{(1 \times 1) + (1 \times 1/2 + 1/2 \times 1/2)}{2} = 7/8.$$

FLGI Probabilities: Computation & Properties

A Non-myopic Group Randomised Procedure

C Just as for the MABP, the computational cost of the exact FLGI probabilities grows with the number of arms (K) and b (block size).

Computation in practice can be done via Monte Carlo simulation. Example: $\boldsymbol{P} = [1\ 1\ ;\ 2\ 1\ ;\ 1\ 2\ ;\ 2\ 2]$ ($K = 4$) and block $b = 9$ then $\boldsymbol{\pi} \approx [0.2646\ ;\ 0.5901\ ;\ 0.0246\ ;\ 0.1208]$ after $5 * 10^2$ replicas.

P1 For equal priors the algorithm defines equal allocation probabilities.

P2 As the block size tends to grow (in the limit it equals the trial size), the design tends to a balanced design (given initial equipoise).

P3 If the block is of only 1 patient (i.e. there is an interim after every patient), the FLGI rule recovers the GI rule.

Introducing Covariates to the Gittins Index Rule

Deriving a CARA FLGI Rule

(1) We consider a MABP with K experimental arms, a control arm and T patients. Before arm k is allocated to patient t, a binary covariate Z_t is observed. Immediately after, a binary response $Y_{t,n}$ is observed.

(2) Reformulate the above MABP: for every treatment-covariate combination there exists a combination arm kz. E.g., the arm "00" corresponds to the control arm and covariate negative patients.

New reformulated MABP has $2(K+1)$ combinations arms (with rate p_{kt}) and patients are optimally allocated to arms with the constraint that they are only allowed arms feasible given their biomarker profile.

(3) We defined a modified GI rule: each patient gets the treatment with the highest GI among the arms available for their biomarker profile.

(4) From this modified GI, a randomised group allocation procedure is defined as in Villar et al (2015) but for every covariate value (and block) we have a different vector of allocation probabilities $\pi_{k,j}(Z)$.

Outline

The CARA FLGI in Practice

Simulation Results

3-arm trial 300 patients $p_{k0} = (0.22; 0.34; \mathbf{0.49})$, $p_{k,1} = (0.47; \mathbf{0.71}; 0.37)$. Treatment-covariate interaction: best arm for covariate negative patients is arm 2 while for covariate positive patients is arm 1.

	Power		Patient Benefit		
	$(1-\beta_0)$	$(1-\beta_1)$	p_0^* (s.d)	p_1^* (s.d)	ENS (s.d)
ER $_{(b=300)}$	0.77	0.69	0.33 (0.04)	0.33 (0.04)	132.01 (8.4)
Thompson $_{(b=10)}$	0.71	0.75	0.49 (0.09)	0.59 (0.09)	149.28 (9.5)
CARA C FLGI $_{(b=10)}$	**0.82**	**0.88**	**0.52 (0.13)**	**0.64 (0.05)**	**151.44 (9.4)**
CARA FLGI $_{(b=10)}$	0.24	0.06	0.78 (0.18)	0.89 (0.13)	171.25 (11.2)
CARA GI $_{(b=1)}$	0.08	0.04	0.70 (0.29)	0.94 (0.06)	171.42 (11.5)

CARA FLGI probabilities (Monte Carlo simulation), $T = 300$, $p_z = 0.5$ and 5000 runs.

- Treatment-covariate interactions are detected by the CARA (Covariate-Adjusted Response Adaptive) FLGI procedure but its statistical power is very low.
- In a multi-armed case the CARA CFLGI addresses the power limitation (though in a two-arm setting power may be insufficient).

Outline

Discussion & Extensions

Trials for Personalised Treatment & Including Patient Benefit Goals

- Derivation of the CARA FLGI probabilities is considerably more complex than competing methods yet its implementation is as feasible as other methods.
- Non-myopic methods offer increased patient benefit (quickly identify best arm) yet have lower power. Power may be improved in a two-armed context at the expense of patient benefit gains.
- Multi-armed trials simulations: patient benefit increases with the number of arms and a protected method achieves both power-ethical advantages ("protecting allocation to control arm" - CARA CFLGI).
- Extensions: polychotomous (every treatment-covariate combination as an arm) and multiple covariates (redefine all the covariates into a single covariate with multiple stratification levels as an arm).
- Fully Bayesian CARA FLGI: larger patient benefit/efficiency gains if historical data is combined with the RAR through the initial prior.

References I

Questions & Comments

Pocock, S. J. and Simon, R. (1975) Sequential treatment assigment with balancing for prognostic factors in the controlled clinical trial *Biometrics*, **31**:103–115.

Taves, D. R., (1974) Minimisation: a new method of assigning patients to treatment and control groups. *Clinical Pharmacology and Therapeutics*, **15**: 443–453.

Wei, L.J., (1978) An application of an urn miodel to the design of sequential controlled clinical trials *Journal of the American Statistical Association*, **73**: 559–563.

Rosenberger, W. F. and Sverdlov, O. (2008). Handling covariates in the design of clinical trials. *Statist. Sci.* **23** 404–419.

Rosenberger, W. F., Vidyashankar, A. N. and Agarwal, D. K. (2001). Covariate-adjusted response-adaptive designs for binary response *J. Biopharm. Statist.* **11** 227–236.

Villar, S., Bowden, J. and Wason, J. (2015) Multi-armed Bandit Models for the Optimal Design of Clinical Trials: Benefits and Challenges. *Statistical Science* **30**, *No. 2, 199-215*.

References II

Questions & Comments

Villar, S. and Rosenberger, W.F. (2018) Covariate-Adjusted Response-Adaptive Randomization for Multi-arm Clinical Trials Using a Modified Forward Looking Gittins Index Rule *Biometrics*. 74(1):49-57.

Villar, S., Wason, J. and Bowden, J. (2015) Response-adaptive Randomization for Multi-arm Clinical Trials using the Forward Looking Gittins Index rule *Biometrics*. 71(4) 969-978

Thompson, W. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294.

John C Gittins and David M Jones (1979) A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, **66**(3):561–565, 1979.