

Asymptotic Properties of Sequential Maximum Likelihood Designs¹

Proprietà asintotiche dei disegni sequenziali di massima verosimiglianza

Alessandro Baldi Antognini

Dipartimento di Scienze Statistiche “Paolo Fortunati”, Università di Bologna
Via Belle Arti 41, 40126, Bologna, Italy; baldi@stat.unibo.it

Riassunto: Nell’ambito degli esperimenti sequenziali per il confronto di due o più trattamenti, molteplici possono essere gli obiettivi sperimentali così come le proporzioni target di assegnazioni dei trattamenti che da essi possono derivare. Il target prefissato spesso dipende dai parametri incogniti del modello ipotizzato. In generale dimostriamo che la procedura sequenziale di massima verosimiglianza, che consiste nell’assegnare ad ogni passo ciascun trattamento con probabilità pari alla stima corrente di massima verosimiglianza della propria proporzione target, converge quasi certamente al target stesso. Inoltre, applicando questa procedura gli stimatori di massima verosimiglianza sono ancora consistenti ed asintoticamente normali e se ne deriva la matrice di varianza-covarianza asintotica.

Keywords: Inference conditional on the design. Sequential design. Response-adaptive randomization. Maximum likelihood.

1. Introduction

Sequential experiments are especially used and indeed often recommended in biomedical practice, but they are also used in other applied contexts. These procedures are very flexible since the experimenter can modify the trial as it goes along, on the basis of the previous allocations and/or observations. In clinical or pharmaceutical trials there is also a need for a certain amount of randomization in the assignments to protect against several types of bias. Furthermore, the experimenter may have several purposes in mind: accurate inference about unknown parameters of the population under investigation, cost saving, ethical preoccupations, etc. Often these objectives can be defined as an optimization problem and sometimes the solutions depend on the unknown parameters of the trial. For example, in the case of a binary response trial for comparing two treatments, let p_1 (q_1) and p_2 (q_2) be the probabilities of success (failure) of T_1 and T_2 ; if we are interested only in inferring on the difference $p_1 - p_2$, it is well known that the so-called Neyman allocation $\frac{\sqrt{p_1 q_1}}{\sqrt{p_1 q_1} + \sqrt{p_2 q_2}}$ is the proportion of assignments of T_1 which minimizes the variance of the estimator $\hat{p}_1 - \hat{p}_2$ and maximizes the power of the corresponding test. However, this allocation does not take into account the patient’s experience and may be unsuitable when ethical issues are involved. In this case one would possibly prefer the allocation $\frac{\sqrt{p_1}}{\sqrt{p_1} + \sqrt{p_2}}$, which minimizes - for a fixed variance of $\hat{p}_1 - \hat{p}_2$ - the expected number of treatment failures. Other “ethical” procedures are the *randomized play-the-winner rule* (Wei and

¹The author was supported by the 2003 Research Project: “Asymptotic properties of stochastic processes: methodological tools and biomedical physical and engineering applications” of the Department of Statistical Sciences, University of Bologna.

Durham, 1978) and the *drop-the-loser rule* (Ivanova, 2003). These designs lead to the very attractive proportion of allocations $q_2/(q_1 + q_2)$, which represents the relative risk of the treatment T_1 (see Hu and Rosenberger, 2003).

Still on the question of comparing two treatments, the problem of finding a response-adaptive design which is asymptotically optimal (in Kiefer's sense) for parametric inference has recently been investigated in Baldi Antognini and Giovagnoli (2004).

In the large majority of the cases, the target proportion of allocations of the treatments depends on the unknown population parameters. A natural response-adaptive randomization procedure consists, at each step, to assign each treatment with probability given by a suitable estimate of the target. But this poses problems as regards the correct inferential paradigm. If the design does not depend on the observed data, then it can be regarded as an ancillary statistics and we may infer conditionally on the design. In contrast, for all response-adaptive procedures, the design variability must be accounted for and we should argue unconditionally (see Rosenberger and Lachin, 2002).

As regards experiments for comparing v treatments, we consider the maximum likelihood design, namely that based on the step-by-step updating of the target estimates by ML. We show that the maximum likelihood design converges to any fixed target proportion of allocations to the treatments. Under this procedure the MLE's retain the strong consistency and asymptotical normality properties and therefore the asymptotic variance-covariance matrix is derived.

2. Sequential designs and likelihood-based inference

Let δ_i be a design point and Z_i the corresponding response. A randomized sequential design can be considered as a stochastic process $\{\delta_i\}_{i \in \mathbb{N}}$ where each δ_i is a random variable. After n steps let $\delta^{(n)} = (\delta_1, \delta_2, \dots, \delta_n)$ be the allocation sequence and $Z^{(n)} = (Z_1, Z_2, \dots, Z_n)$ the sequence of the corresponding responses; the choice of the next design point δ_{n+1} will depend on the previous allocations and/or outcomes $(\delta^{(n)}, Z^{(n)})$, so that the next allocation can be represented by the conditional distributions $\mathfrak{R}(\delta_{n+1} | \delta^{(n)}, Z^{(n)})$. Let $\mathcal{L}_i$ be the probability law of the response Z_i . In what follows we assume that

- $\mathcal{L}_i$ depends on the design point δ_i where it is observed and a vector of unknown parameters and satisfies the usual Rao-Cramer regularity conditions;
- conditionally on the design points, the observations are independent.

Thus, the conditional distribution $\mathfrak{R}$ of the design point is ancillary for the likelihood of the parameters, whether or not the design is adaptive, and the maximum likelihood estimators are the same as if the observed design points were predetermined. However, in general, the distribution of the MLE's is not the same as when the observations are *i.i.d.*, so that confidence intervals and hypothesis testing cannot be constructed in the classical way. Chaudhuri and Mykland (1993), Rosenberger, Flournoy and Durham (1997) and, for the case of two treatments, Melfi and Page (2000) give some regularity and convergence conditions under which the MLE's retain the strong consistency and asymptotical normality properties, but the problem lies in the possibility of checking if such conditions hold for a given application (Rosenberger and Lachin, 2002).

3. Designs for comparing v treatments: notation and model

Suppose we wish to compare v treatments $T_1, \dots, T_v$ and let $\delta_k = (\delta_{1k}, \dots, \delta_{vk})$ be the treatment allocation at step k , where $\delta_{jk} = 1$ if T_j is assigned to the k -th unit, 0 otherwise. The experimental design can be represented by the law of δ_1 and the sequence $\{Pr(\delta_{j(n+1)} = 1 | \delta^{(n)}, Z^{(n)}), j = 1, \dots, v\}_{n \in \mathbb{N}}$. In addition, let $N_j(k) = \sum_{i=1}^k \delta_{ji}$ and $\pi_{jk} = k^{-1} N_j(k)$ be the number and the proportion of allocations to T_j ($j = 1, \dots, v$) respectively, and $X_{1k}, \dots, X_{vk}$ be the potential responses to each treatment. At any step k , exactly one of them is observed and the response Z_k can be written as $Z_k = \sum_{j=1}^v \delta_{jk} X_{jk}$. For any given $j = 1, \dots, v$ we assume that $\{X_{jk}\}_{k \in \mathbb{N}}$ are *i.i.d.* and the distributions belong to the same regular exponential family, i.e.

$$X_j \sim \exp \left\{ \phi_j^{-1} [a(\theta_j)x_j - b(\theta_j)] + c(x_j, \phi_j) \right\} \quad (1)$$

with $a(\cdot)$ and $b(\cdot)$ twice continuously differentiable, where $\theta_j \in \Theta \subseteq \mathbb{R}$ is the parameter of interest and $\phi_j \in \mathbb{R}^+$ is a nuisance parameter. Special cases are normal, Poisson, exponential, Bernoulli and logit models for the responses.

4. Maximum likelihood designs

Suppose now there exists a given target proportion $\rho = (\rho_1, \dots, \rho_v)$, with $\sum_{j=1}^v \rho_j = 1$, of allocations of the v treatments - which is considered optimal by the experimenter - depending on the unknown population parameters $\theta = (\theta_1, \dots, \theta_v)$ and $\phi = (\phi_1, \dots, \phi_v)$. In particular we assume that ρ is a continuous function of $(\theta; \phi)$ with

$$\rho_j = \rho_j(\theta; \phi) > 0 \quad \forall j = 1, \dots, v. \quad (2)$$

Classical examples are A-, D- and E-optimal proportions for parameter estimation and the target allocations corresponding to the other “ethical” procedures considered previously. For hypothesis testing, other proportions would be suitable (Hu and Rosenberger, 2003). Additionally, suppose that an initial number $(n_{10}, \dots, n_{v0})$ of observations have been allocated to the treatments in order to compute an initial non-trivial estimate (Melfi, Page and Geraldes, 2001) and at any stage $k \geq n_0 = \sum_{j=1}^v n_{j0}$ let the MLE’s $\hat{\theta}_k = (\hat{\theta}_{1k}, \dots, \hat{\theta}_{vk})$ and $\hat{\phi}_k = (\hat{\phi}_{1k}, \dots, \hat{\phi}_{vk})$ of θ and ϕ based on $(\delta^{(k)}, Z^{(k)})$. Then we define the maximum likelihood design by choosing at each step treatment T_j with probability given by the current estimate of the corresponding target proportion ρ_j , namely

$$Pr(\delta_{j(k+1)} = 1 | \delta^{(k)}, Z^{(k)}) = \hat{\rho}_{jk} = \rho_j(\hat{\theta}_k; \hat{\phi}_k) \quad \forall k \geq n_0. \quad (3)$$

Theorem 1 *For any given target ρ satisfying (2), under the maximum likelihood design $N_j(n) \rightarrow \infty$ almost surely as $n \rightarrow \infty$ for any $j = 1, \dots, v$. Thus, the MLE’s of θ , ϕ and ρ are strongly consistent and*

$$\lim_{n \rightarrow \infty} N_j(n) / n = \rho_j \quad a.s. \quad \forall j = 1, \dots, v.$$

The proof of this result generalizes Theorem 8 and Corollary 9 of Baldi Antognini and Giovagnoli (2004) to the case of $v \geq 2$ treatments and any given target.

Without loss of generality, the exponential family in (1) can always be reparametrized

in such a way that the parameters of interest are the mean responses of the treatments, namely $E[X_j] = \theta_j$, so that for any $j = 1, \dots, v$ the MLE's of θ_j are the sample means $\hat{\theta}_{jn} = N_j(n)^{-1} \sum_{i=1}^n \delta_{ji} X_{ji}$. Given this reparametrization there follows

Theorem 2 Let $Var(X_j) = \sigma_j^2$; then for any given target ρ satisfying (2), under the maximum likelihood design the vector $(\sqrt{n} [\hat{\theta}_{jn} - \theta_j], j = 1, \dots, v)$ is asymptotically normal with mean zero and diagonal variance-covariance matrix with elements $\nu_{jj} = \sigma_j^2 / \rho_j$ for $j = 1, \dots, v$ ($\nu_{jl} = 0 \forall j \neq l$).

Example Given v treatments, assume $X_j \sim N(\mu_j; \sigma_j^2)$ for any $j = 1, \dots, v$. The A-optimal proportions $\rho_j = \sigma_j / \sum_{i=1}^v \sigma_i$ minimizes the sum of the variances of the estimators of the means. At step k , let $\hat{\mu}_{jk}$ and $\hat{\sigma}_{jk}$ be the MLE's of μ_j and σ_j ; by (3) the maximum likelihood design is

$$\Pr(\delta_{jk+1} = 1 | \delta^{(k)}, Z^{(k)}) = \hat{\sigma}_{jk} / \sum_{i=1}^v \hat{\sigma}_{ik}, \quad j = 1, \dots, v \quad \forall k \geq n_0.$$

By Theorem 1 such procedure converges almost surely to the target, namely

$$\lim_{n \rightarrow \infty} N_j(n) / n = \sigma_j / \sum_{i=1}^v \sigma_i \quad a.s. \quad \forall j = 1, \dots, v$$

and by Theorem 2, as $n \rightarrow \infty$

- $\sqrt{n} [\hat{\mu}_{jn} - \mu_j] \rightarrow N(0, \sigma_j \sum_{i=1}^v \sigma_i) \quad j = 1, \dots, v;$
- $\sqrt{n} [(\hat{\mu}_{jn} - \hat{\mu}_{ln}) - (\mu_j - \mu_l)] \rightarrow N(0, (\sigma_j + \sigma_l) \sum_{i=1}^v \sigma_i) \quad \forall j, l \quad j \neq l.$

References

- Baldi Antognini, A. and Giovagnoli, A. (2004) On the large sample optimality of sequential designs for comparing two treatments. *Manuscript submitted*.
- Chaudhuri, P. and Mykland, P.A. (1993) Nonlinear experiments: optimal design and inference based on likelihood. *Journal of the American Statistical Association*, **88**, 538-46.
- Hu, F. and Rosenberger, W.F. (2003) Optimality, variability, power: evaluating response-adaptive randomization procedures for treatment comparisons. *Journal of the American Statistical Association*, **98**, 671-78.
- Ivanova A. (2003) A play-the-winner-type urn design with reduced variability. *Metrika*, **58**, 1-13.
- Melfi, V. and Page, C. (2000) Estimation after adaptive allocation. *Journal of Statistical Planning and Inference*, **87**, 353-63.
- Melfi, V., Page, C. and Geraldles, M. (2001) An adaptive randomized design with application to estimation. *Canadian Journal of Statistics*, **29**, 107-16.
- Rosenberger, W.F., Flournoy, N. and Durham, S.D. (1997) Asymptotic normality of maximum likelihood estimators from multiparameter response-driven designs. *Journal of Statistical Planning and Inference*, **60**, 69-76.
- Rosenberger, W.F. and Lachin, J.M. (2002) *Randomization in Clinical Trials*. Wiley.
- Wei, L.J. and Durham, S.D. (1978) The randomized play-the-winner rule in medical trials. *Journal of the American Statistical Association*, **73**, 840-43.