THE UNIVERSITY OF
WARWICK

Original citation:
Rossell, David and Muller, P.. (2013) Sequential stopping for high-throughput experiments. Biostatistics, Volume 14 (Number 1). pp. 75-86. ISSN 1465-4644

Permanent WRAP url:
http://wrap.warwick.ac.uk/53407

Copyright and reuse:
The Warwick Research Archive Portal (WRAP) makes this work of researchers of the University of Warwick available open access under the following conditions.

This article is made available under the Creative Commons Attribution 3.0 Unported (CC BY 3.0) license and may be reused according to the conditions of the license. For more details see: http://creativecommons.org/licenses/by/3.0/

A note on versions:
The version presented in WRAP is the published version, or, version of record, and may be cited as it appears here.

For more information, please contact the WRAP Team at: publications@warwick.ac.uk

warwick**publications**wrap

highlight your research

http://wrap.warwick.ac.uk

Sequential stopping for high-throughput experiments

DAVID ROSSELL*

Biostatistics and Bioinformatics Unit, Institute for Research in Biomedicine of Barcelona, Barcelona 08028, Spain

david.rossell@irbbarcelona.org

PETER MÜLLER

Department of Mathematics, UT Austin, Austin 78712, TX, USA

Summary

In high-throughput experiments, the sample size is typically chosen informally. Most formal sample-size calculations depend critically on prior knowledge. We propose a sequential strategy that, by updating knowledge when new data are available, depends less critically on prior assumptions. Experiments are stopped or continued based on the potential benefits in obtaining additional data. The underlying decision-theoretic framework guarantees the design to proceed in a coherent fashion. We propose intuitively appealing, easy-to-implement utility functions. As in most sequential design problems, an exact solution is prohibitive. We propose a simulation-based approximation that uses decision boundaries. We apply the method to RNA-seq, microarray, and reverse-phase protein array studies and show its potential advantages. The approach has been added to the Bioconductor package `gaga`.

Keywords: Decision theory; Forward simulation; High-throughput experiments; Multiple testing; Optimal design; Sample size; Sequential design.

1. Introduction

In high-throughput studies (HTSs), the sample size is usually chosen informally. The resulting experiment may either be not informative enough or unnecessarily extensive. To address this problem, we develop a sequential design framework for HTSs. That is, we investigate the question whether the currently available data in a typical HTS suffices and, if not, how to determine the optimal stopping strategy. We focus on experiments to perform group comparisons, although our ideas remain useful for other inferential goals. For simplicity, we discuss the two-group case, but our software allows >2 groups. The proposal is based on Bayesian decision theory, so that decisions are coherent with respect to an underlying utility function and probability model. We emphasize ease of interpretation and use.

Several authors proposed fixed sample-size calculations for HTSs, i.e. the sample size is fixed at the beginning of the experiment (Dobbin and Simon, 2007; Lee and Whitmore, 2004; Müller *and others*, 2004; Pan *and others*, 2002; Zien *and others*, 2003). The main limitations are the lack of flexibility to

---

*To whom correspondence should be addressed.

incorporate new data and the need for a good prior guess of certain features, e.g. effect sizes or the proportion of differentially expressed (DE) genes.

In contrast, sequential sample-size designs update knowledge and make decisions as data are collected, i.e. they are robust with respect to prior choices. A sequential design stops or continues experimentation on the basis of all available data. A potential drawback is the need to carry out experimentation in batches. The associated increase in time or experimentation costs may outweigh the potential advantages. This should not be a major concern in many HTSs, as most high-throughput technologies (e.g. microarrays, sequencing, or mass spectrometry) process samples in small batches. Assessing the promise of continuing experimentation after each batch seems natural. Also, samples may be costly to obtain or there may be ethical concerns, e.g. in human studies. These situations offer great potential for sequential strategies.

Ruppert *and others* (2007), Tibshirani (2006), and Ferreira and Zwinderman (2006) proposed two-step designs focused on microarray differential expression problems. Two-step designs adapt to the observed data to a limited extent. Gibbons *and others* (2005) and Durrieu and Briollais (2009) propose sequential designs that select a single/few genes and stop the trial when differences in expression can be estimated with high precision. The focus on a few genes limits the application to HTS. Researchers typically use HTS as a screening test to identify candidates, which are then validated with more precise techniques (e.g. real-time PCR). The usual goal is not to estimate differential expression accurately but to find promising targets. The Durrieu and Briollais (2009) model is appropriate for paired observations, e.g. two-channel arrays.

We propose an approach for unpaired data that screens a large number of candidates and attempts to maximize the number of promising targets. The framework is directly applicable to many probability models and experiments, including sequencing, microarrays, and reverse-phase protein arrays (RPPAs). With minor modifications, it can be adapted to other experimental goals. The main hurdle with decision-theoretic optimal sequential designs is the prohibitive computational cost, even in single-outcome experiments. Rossell *and others* (2006) developed an approach based on the ideas of Müller *and others* (2006), Brockwell and Kadane (2003), and Carlin *and others* (1998). They compute a summary statistic S each time new data are observed and they use decision boundaries that partition the sample space. The experiment is terminated when S first falls in the stopping region. The sequential problem is reduced to the (non-sequential) problem of finding optimal boundaries. The choice of these boundaries accounts for all future data, which distinguishes the solution from myopic approximations. Here we extend these ideas to high-dimensional data and apply them to differential expression problems.

Section 2 formalizes the problem and two convenient probability models. Section 3 describes sequential stopping and the infeasibility of an exact solution. Section 4 proposes an approximate solution. Section 5 presents examples and Section 6 some concluding remarks. The supplementary material available at *Biostatistics* online contains theoretical and practical considerations, and an example with the R code.

2. Data format and model

We motivate the discussion in the context of experiments that study differential gene expression, but the proposal remains applicable to other setups. Let n be the number of outcomes (e.g. genes) and T be the maximum sample size. T is usually determined by budget constraints, accrual rates, or an informed guess. Let x_{ij} be the measurement for gene $i = 1, \ldots, n$ and sample $j = 1, \ldots, T$, and $z_j \in \{0, \ldots, n_z\}$ be the group of sample j. For simplicity, here we assume $z_j \in \{0, 1\}$, i.e. we compare two groups. Generalization to $n_z > 1$ is straightforward and is implemented in the `gaga` package.

A latent variable $\delta_i = 1$ indicates that gene i is DE across groups and $\delta_i = 0$ that it is equally expressed (EE). The indicator δ_i represents the unknown truth and is part of the parameter vector. Let $\boldsymbol{\theta}_i$ be parameters indexing a probability model for $(x_{i1}, \ldots, x_{iT})$. Optionally, let $\boldsymbol{\omega}$ be additional hyper-parameters. For example, $\boldsymbol{\omega}$ could index a regression on important covariates. Let $\mathbf{x}_t = \{x_{it}, 1 \leqslant i \leqslant n\}$ be the data obtained

at time t and $\mathbf{x}_{1:t} = \{x_{ij}, 1 \leqslant i \leqslant n, 1 \leqslant j \leqslant t\}$ be all data available up to time t. Further, let $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_n)$ and $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_n)$.

2.1 *Probability model*

Our proposal requires extensive predictive simulation and model fitting. Hence, the model must be computationally efficient. For instance, the examples in Section 5 required posterior inference in millions of simulated datasets. On the other hand, the model needs to be sufficiently flexible to capture the important features of the data.

Here we use the GaGa (Rossell, 2009) and log-normal normal with generalized variance (NN) models (Yuan and Kendziorski, 2006) models to illustrate the approach. Both offer a reasonable compromise between flexibility and computational cost. The GaGa model assumes $x_{ij} \sim \text{Ga}(\alpha_i, \alpha_i/\lambda_{iz_j})$. The NN model uses $x_{ij} \sim N(\mu_{iz_j}, \sigma_i^2)$. The triple $\boldsymbol{\theta}_i = (\lambda_{i0}, \lambda_{i1}, \alpha_i)$ (GaGa) or $\boldsymbol{\theta}_i = (\mu_{i0}, \mu_{i1}, \sigma_i^2)$ (in the NN model) incorporates gene-specific variability and gene-by-group specific means. A hierarchical prior on $\boldsymbol{\theta}_i$ assigns positive prior probability to means being equal across groups. The GaGa hierarchical prior is

$$\lambda_{i0}^{-1} \sim \text{Ga}(\alpha_0, \alpha_0/\nu), \quad \alpha_i | \beta, \mu \sim \text{Ga}(\beta, \beta/\mu),$$

$$\lambda_{i1}^{-1} \mid \lambda_{i0}, \delta_i \sim \begin{cases} \text{Ga}(\alpha_0, \alpha_0/\nu) & \text{if } \delta_i = 1, \\ I(\lambda_{i1} = \lambda_{i0}) & \text{if } \delta_i = 0, \end{cases} \tag{2.1}$$

and $P(\delta_i = 1) = \pi$, independently across i. The hyper-parameters are $\boldsymbol{\omega} = (\alpha_0, \nu, \beta, \mu, \pi)$. The NN hierarchical prior is $\mu_{i0} \sim N(\mu_0, \tau_0^2)$, $\sigma_i^{-2} \sim \text{Ga}(\nu_0/2, \nu_0\sigma_0^2/2)$, also with probability of ties $P(\delta_i = 1) = \pi$, independently across i. For this model, $\boldsymbol{\omega} = (\mu_0, \tau_0, \nu_0, \sigma_0, \pi)$. The GaGa sampling distribution for x_{ij} captures asymmetries that are frequently observed in HTS. The NN assumptions are similar to those of the popular limma approach (Smyth, 2004), which has been found useful in many applications. The supplementary material available at *Biostatistics* online (Section 3) shows goodness-of-fit assessments that help choose the most appropriate model for a particular dataset.

In terms of computational complexity, conditional on $\boldsymbol{\omega}$ the posterior distributions are available in closed form. We treat $\boldsymbol{\omega}$ as fixed, avoiding the need for Markov Chain Monte Carlo (MCMC) simulation. This substantially increases the computational speed. We estimate $\boldsymbol{\omega}$ via expectation-maximization as in Rossell (2009) and Yuan and Kendziorski (2006). The latter proposed a method of moments estimate for (ν_0, σ_0^2), which can result in overestimating π. We illustrate this issue and outline a simple procedure to adjust $\hat{\pi}$ in the supplementary material available at *Biostatistics* online (Section 4).

While we use these two models in our examples, the upcoming discussion of the optimal stopping policy remains valid for any alternative probability model.

2.2 *Pre-processing*

We assume that the data are suitably pre-processed. This is critical for meaningful inference. For instance, ignoring batch effects may bias or add uncertainty to group comparisons. We note that some technologies such as RNA-seq may be less sensitive to batch effects, and that these can be partially mitigated by good design, e.g. by balancing the number of samples in each group and batch. We recommend jointly pre-processing data after every batch, as some technical biases (e.g. probe or GC-content biases) may be better assessed once more data are collected.

Batch effects and other sources of variability may be either addressed in the pre-processing or in the analysis by including appropriate terms in the model. Following Yang and Speed (2002) and Durrieu and Briollais (2009), we argue in favor of the former. As an illustration, let $\mathbf{y}_{ij}$ be a vector of

covariates that are used in the adjustment, and assume that $E(x_{ij}|\mu_{iz_j}, \mathbf{y}_{ij}) = \mu_{iz_j} + g(\mathbf{y}_{ij})$. Here, $g(\cdot)$ captures the effect of $\mathbf{y}_{ij}$ on the outcome, and could represent a non-linear adjustment that cannot be captured by the analysis model. One could then use the partial residuals $\tilde{x}_{ij} = x_{ij} - \hat{g}(\mathbf{y_{ij}})$ as the pre-processed data, where $\hat{g}(\cdot)$ is an appropriate estimate of $g(\cdot)$.

We note that the domain of the data must match the assumptions of the model. For example, while most technologies deliver positive expression measurements, pre-processed data may sometimes present negative values that are not allowed in the GaGa model. A simple strategy to deal with negative values is to define $\tilde{x}_{iz_j} = x_{iz_j} + k$, where the offset $k > 0$ ensures that $\tilde{x}_{iz_j} > 0$. Alternatively, define $\tilde{x}_{iz_j} = e^{x_{iz_j}}$, but this option may produce outliers that decrease the model goodness-of-fit. In practice, we recommend trying several transformations and producing some goodness-of-fit plots (e.g. see supplementary material available at *Biostatistics* online, Section 3).

3. OPTIMAL SEQUENTIAL STOPPING

3.1 *Decision criterion*

We formalize sequential sample-size calculation within a Bayesian decision-theoretic framework. The optimal design is chosen by maximizing the expectation of an appropriate utility function. At each decision time, the expected utility is conditional on all available data (which may be no data at all) and averaged with respect to uncertainty on the model parameters and future data, assuming optimal future decisions.

It is convenient to distinguish sequential and terminal decisions. Sequential decisions correspond to stopping vs. continuation and are made after each batch of observations. Terminal decisions are the classification of genes into EE ($\delta_i = 0$) or DE ($\delta_i = 1$), and are taken only upon stopping. Let $s_t = s(\mathbf{x}_{1:t}) = 1$ indicate the sequential decision of stopping at time t and let $s_t = 0$ indicate continuation. Equivalently, s_t can be described by the stopping time $\tau = \min\{t : s_t = 1\}$. We use s_t and τ interchangeably. Let $d_i(\mathbf{x}_{1:t}) = 1$ (0) indicate the terminal decision to report gene i as DE (EE). Also, let $\mathbf{d}(\mathbf{x}_{1:t}) = (d_1(\mathbf{x}_{1:t}), \ldots, d_n(\mathbf{x}_{1:t}))$. Both s_t and $\mathbf{d}(\mathbf{x}_{1:t})$ depend on all data available up to time t.

In a fully decision-theoretic approach sequential and terminal decisions are chosen to jointly maximize the expected utility. Instead, we assume a fixed rule for $\mathbf{d}(\mathbf{x}_{1:t})$ and focus on the optimal choice of s_t only. We take terminal decisions using the Bayes rule of Müller *and others* (2004) to control the posterior expected false discovery rate (FDR) below some specified level. The posterior expected FDR is $(1/D) \sum_i d_i(\boldsymbol{x}_{1:t})[1 - E(\delta_i \mid \boldsymbol{x}_{1:t})]$, where $D = \sum_i d_i(\boldsymbol{x}_{1:t})$ is the number of reported positives. We use the 0.05 level throughout.

Sequential stopping decisions s_t are based on a utility function with sampling cost c and a unit reward for each correctly identified DE outcome

$$u(s_t = 1, \mathbf{d}(\boldsymbol{x}_{1:\tau}), \boldsymbol{x}_{1:\tau}, \boldsymbol{\delta}) = -c\tau + \sum_{i=1}^{n} \delta_i d_i(\boldsymbol{x}_{1:\tau}). \tag{3.1}$$

The second term in (3.1) is the number of true positives (TPs). The cost c is the minimum number of TPs that make it worthwhile to obtain one more sample. This interpretation allows for easy elicitation of c, without any reference to the formal mathematical framework. The utility function (3.1) focuses on statistical rather than biological significance, as the size of the effect is not considered. A simple alternative is obtained by substituting $|\mu_{i1} - \mu_{i2}|\delta_i d_i(\boldsymbol{x}_{1:\tau})$ in the summation in (3.1); see Müller *and others* (2004) or Rice *and others* (2008) for other interesting alternatives. The upcoming discussion is independent of the specified utility.

3.2 *Optimal rule*

The optimal stopping decision s_t maximizes $u(\cdot)$, in expectation over all unknowns, including parameters $(\boldsymbol{\theta}, \boldsymbol{\delta})$, and future data $\mathbf{x}_{\tau+1:T}$. An exact solution requires dynamic programming, also known as backward induction (DeGroot, 2004). At time t, the optimal decision is to stop if the posterior expected utility for $s_t = 1$, denoted by $\bar{u}_t(s_t = 1, \mathbf{x}_{1:t})$, is greater than $\bar{u}_t(s_t = 0, \mathbf{x}_{1:t})$. Evaluating $\bar{u}_t(s_t = 1, \mathbf{x}_{1:t})$ is usually straightforward. For (3.1), we find

$$\bar{u}_t(s_t = 1, \boldsymbol{x}_{1:t}) = -ct + \sum_{i=1}^{n} P(\delta_i = 1 \mid \boldsymbol{x}_{1:t}) d_i(\boldsymbol{x}_{1:t}), \tag{3.2}$$

where $P(\delta_i = 1 | \boldsymbol{x}_{1:t})$ is the posterior probability that outcome i is DE. The expectation is with respect to $\boldsymbol{\delta}$ only, as we fix the terminal decision $\mathbf{d}(\mathbf{x}_{1:t})$. The posterior probability $P(\delta_i = 1 | \mathbf{x}_{1:t})$ can be computed in closed form for some models (including GaGa and NN) or can easily be estimated from the MCMC output. Evaluating $\bar{u}_t(s_t = 0, \boldsymbol{x}_{1:t})$ is more challenging. An exact solution requires assessing the expected utility for all possible future data trajectories $\boldsymbol{x}_{t+1}, \ldots, \boldsymbol{x}_T$, substituting the optimal decisions $s_{t+1}, \ldots, s_T$. The computational cost is prohibitive.

4. Approximation by optimal decision boundaries

Berry *and others* (2001), Brockwell and Kadane (2003), DeGroot (2004), and Müller *and others* (2006) discuss alternatives to an exact optimal sequential solution. Following Rossell *and others* (2006), we define sequential stopping boundaries. We restrict the maximization to rules that depend on the data $\boldsymbol{x}_{1:t}$ only through a summary statistic S_t and linear boundaries that partition the sample space. We propose using $S_t = \Delta_t U$, where

$$\Delta_t U \equiv E_{\boldsymbol{x}_{t+1}}[\bar{u}_{t+1}(s_{t+1} = 1, \mathbf{x}_{1:t+1}) \mid \boldsymbol{x}_{1:t}] - \bar{u}_t(s_t = 1, \boldsymbol{x}_{1:t}) - c$$

is the one-step ahead increase in the expected utility and $E_{\boldsymbol{x}_{t+1}}(\cdot | \boldsymbol{x}_{1:t})$ conditions on $\boldsymbol{x}_{1:t}$ and marginalizes with respect to future data $\boldsymbol{x}_{t+1}$. For (3.1), we find $\Delta_t U = \Delta_t(\text{TP}) - c$, i.e. $\Delta_t U$ is the expected increase in the TPs, and decision boundaries can equivalently be written in terms of $\Delta_t(\text{TP})$.

Consider the example in Figure 1. The thick black line is a decision boundary. Every time we observe new data, we compute $\Delta_t(\text{TP})$. If $\Delta_t(\text{TP})$ lies above the boundary, we continue experimentation, otherwise we stop. That is, we experiment as long as enough new TPs are expected.

Let $\boldsymbol{b} = (b_0, b_1)$ be the intercept and slope defining the linear boundaries, and $U(\boldsymbol{b}, \boldsymbol{x}_{1:t})$ be the associated expected utility given data up to time t. In other words, $U(\boldsymbol{b}, \boldsymbol{x}_{1:t})$ is the expected utility conditional on $\boldsymbol{x}_{1:t}$ when the stopping decision is based on a decision boundary indexed by $\boldsymbol{b}$. Algorithm 1 details a forward simulation algorithm (Carlin *and others*, 1998) to evaluate the required expectations, and a grid search to carry out the maximization of $U(\boldsymbol{b}, \boldsymbol{x}_{1:t})$ with respect to $\boldsymbol{b}$. The algorithm assumes that t samples are already available. For no data use $t = 0$ and $\boldsymbol{x}_{1:t} = \emptyset$ (see Section 5.3).

Algorithm 1. Optimal sequential boundary determination.

(1) *Forward simulation*. Simulate $\boldsymbol{x}_{t+1:T}^{(j)}$ from the posterior predictive $P(\boldsymbol{x}_{t+1:T} \mid \boldsymbol{x}_{1:t})$, $j = 1, \ldots, B$. For each $\boldsymbol{x}_k^{(j)}$, compute $\Delta_t(\text{TP})^{(j)}$, $k = t+1, \ldots, T$.

(2) *Grid search*. For each $\boldsymbol{b}$, find the stopping times $\tau^{(j)}$ for all saved trajectories $\Delta_t(\text{TP})^{(j)}$.

(3) *Optimum*. Select $\boldsymbol{b}^\star \equiv \arg\max_{\boldsymbol{b}} \{\bar{U}(\boldsymbol{b}, \boldsymbol{x}_{1:t})\}$, where $\bar{U}(\boldsymbol{b}, \boldsymbol{x}_{1:t}) = (1/B) \sum_{j=1}^{B} \bar{u}(s_{\tau^{(j)}} = 1, \boldsymbol{x}_{1:\tau^{(j)}}^{(j)})$.

Fig. 1. A GaGa model based optimal sequential boundary for $c = 50$ (thick black line) and forward simulation trajectories (light gray lines), for example, in Section 5.1.

Figure 1 shows simulated $\Delta_t(\text{TP})$ as gray lines. For each boundary $\boldsymbol{b}$, we determine the stopping time for each trajectory and average the expected terminal utilities. At $t = T - 1$, we do not determine stopping using $\boldsymbol{b}$ but the optimal rule $\Delta_{T-1}U > 0$.

In principle $\boldsymbol{b}^{\star}$ can be recomputed every time new data are observed. Recomputation can help to decide between multiple optima and update $P(\boldsymbol{x}_{t+1:T} \mid \boldsymbol{x}_{1:t})$. In our examples, we determine $\boldsymbol{b}^{\star}$ only once, either based on a pilot dataset or prior knowledge, but we indicate the usefulness of recomputation when appropriate.

In addition to the intuitive appeal, some theoretical considerations motivate our approach. First, fixed-sample designs are special cases, e.g. $\boldsymbol{b} = c(4.5, \infty)$ results in a fixed sample size of 5. The myopic rule of continuing as long as $\Delta_t U > 0$ (Berry and Fristedt, 1985, Chapter 7) is the special case $\boldsymbol{b} = (c, 0)$. We generalize the idea with an arbitrary boundary on $\Delta_t U$. An important assurance is that $\Delta_t \text{TP}$ converges to 0 as $t \to \infty$, which guarantees eventual stopping; see supplementary material available at *Biostatistics* online (Section 1) for a formal statement and proof.

5. EXAMPLES

We compare our approach and the fixed sample designs of Müller *and others* (2004) in several important experimental conditions. The supplementary material available at *Biostatistics* online discuss pre-processing and goodness-of-fit (Section 3) and an additional RNA-seq example with the R code (Section 4).

5.1 *Simulated microarray study*

We plan collecting data in batches of 2 arrays per group, with a maximum of 20 per group (i.e. T=10 batches). Recall that c is the minimum number of new DE genes that compensate the cost of one more batch. We consider $c = 25$, 50, and 100. To keep the simulation realistic, we estimated the hyper-parameters based on data from a study of leukemia microarray data (Armstrong *and others*, 2002). We focus on 24 acute lymphoblastic leukemia (ALL) and 18 mixed-lineage leukemia (MLL) trans-location samples. The

Table 1. *Simulated data.* t_F^* : *fixed sample size*; t_S^* : *average sequential sample size*; $U_S^* - U_F^*$: *expected utility for sequential design minus the expected utility for a fixed sample*

c	t_F^*	$\pi = 0.5\hat{\pi}$		$\pi = \hat{\pi}$		$\pi = 2\hat{\pi}$	
		t_S^*	$U_S^* - U_F^*$	t_S^*	$U_S^* - U_F^*$	t_S^*	$U_S^* - U_F^*$
GaGa							
25	7	6.0	7.7	7.1	0.4	10.0	34.2
50	5	4.1	16.6	5.0	0.0	6.9	28.7
100	3	3.0	0.0	3.0	0.0	4.0	58.6
NN							
25	7	5.6	11.1	7.2	0	10.0	32.4
50	4	3.2	10.8	4.1	0.1	5.7	51.2
100	3	2.0	41.9	3.0	0.1	3.0	0

estimated proportion of DE genes is $\hat{\pi}$=0.063 under the GaGa model and 0.05 for the NN model. We find optimal strategies based on $\pi = \hat{\pi}$, but we assess performance under model mis-specification by also simulating data under $\pi = 0.5\hat{\pi}$ and $\pi = 2\hat{\pi}$, while leaving $\pi = \hat{\pi}$ unchanged in the analysis model. We obtained 250 simulations under each scenario.

Figure 1 shows the optimal boundary for $c = 50$ and simulated $\Delta_t(\text{TP})$ (gray lines) under the GaGa model. Table 1 reports expected utilities and stopping times. The optimal fixed sample sizes for $c = 50$ under the GaGa and NN models are $t_F^* = 5$ and $t_F^* = 4$, respectively. When $\pi = \hat{\pi}$, the expected sequential sample sizes are 5.0 and 4.1 (respectively) and there is no gain in the posterior expected utility. Sequential designs offer little advantages when the data match the prior expectations. However, when prior expectations are unrealistic, sequential designs adapt to the observed data. When π was overstated by the prior ($\pi = 0.5\hat{\pi}$), sequential designs stopped earlier than the fixed sample-size designs. Conversely, when $\pi = 2\hat{\pi}$, they stopped later so that more DE genes could be found. For instance, for $c = 50$ the GaGa sequential design requires 4.1 data batches when $\pi = 0.5\hat{\pi}$ and 6.9 when $\pi = 2\hat{\pi}$. The fixed design always requires 5.

5.2 *High-throughput sequencing example*

We use a pilot RNA-seq dataset with two muscle and one brain human samples to design two hypothetical studies. Study 1 compares gene expression for muscle vs. brain. Many DE genes are expected. Hypothetical Study 2 compares the two muscle samples. No genes should be DE. In both cases, we use one sample per group as pilot data. We consider up to $T = 5$ more samples, in batches of one sample. The GaGa model provided a reasonable fit to these data (supplementary material available at *Biostatistics* online, Section 3).

We determined the optimal boundary for sampling costs $c = 0, 1, \ldots, 100$. Figure 2(a) shows that $\Delta_t(\text{TP})$ is maximal for $t = 2$ additional data batches. As suggested by Theorem 1 (supplementary material available at *Biostatistics* online, Section 1), the incremental reward decreases as t grows further. The dashed boundary shows that, for $c > 66$, the optimal decision is to stop experimentation. For $c \leqslant 66$, there are multiple optimal b^*. The solid black lines show two optima. In both cases, the decision at $t = 0$ is to continue. Since the simulated trajectories do not cross either boundary, we expect experimentation to continue up to $T = 5$. The future realized $\Delta_t(\text{TP})$ might cross the boundary, in which case the design would adapt and stop experimentation before $T = 5$. Given that the pilot data contains one sample per group, we would re-determine $b^\star$ upon observing new data.
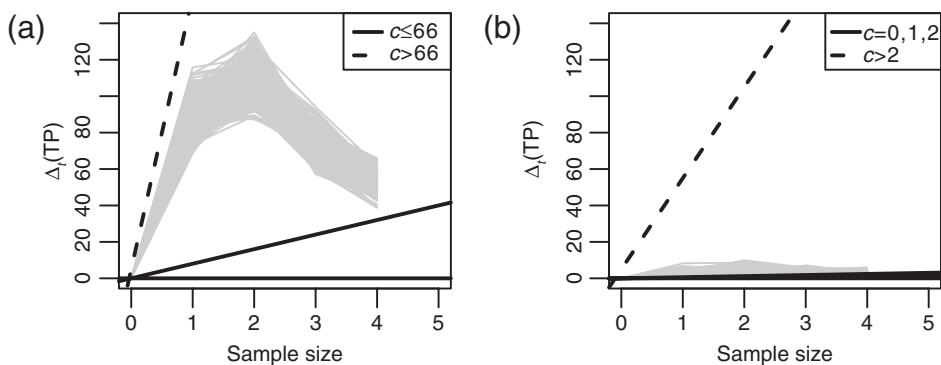
Fig. 2. Simulated one-step expected increase in true discoveries $\Delta_t(\text{TP})$ (gray lines) and optimal boundaries for several sampling costs c (black lines). Left: brain vs. muscle (two multiple optimal boundaries shown for $c \leqslant 66$), right: muscle vs. muscle.

The hypothetical muscle vs. muscle comparison simulation is shown in Figure 2(b). In this case, $\Delta_t(\text{TP})$ is negligible and the optimal design is to stop at $t = 0$ (i.e. not to collect any further data) for any $c > 2$. The result seems sensible as no DE genes are expected.

5.3 *Microarray example*

We consider the leukemia study of Campo Dell'Orto *and others* (2007) recording mRNA expression for 21 ALL and 15 MLL patients and 54 675 genes. We consider designing the study before any data were available. In such circumstances, one could estimate the hyper-parameters $\boldsymbol{\omega}$ from a similar study. We used the Armstrong *and others* (2002) study (Section 5.1) as it was also carried on ALL/MLL patients and used the same microarray platform. Once fixed and sequential designs were determined, we used the historical data to compare performance. We use batches of two arrays per group, maximum $T = 7$ batches, and $c = 50$.

The white bars in Figure 3 (left panels) show the expected utility under the GaGa and the NN priors for all fixed sample sizes. The optimal fixed sample sizes are $t_F^* = 5$ batches (GaGa) and $t_F^* = 4$ (NN). The right panels show the optimal boundaries and $\Delta_t(\text{TP})$ computed from the observed data up to time $t = 1, \ldots, 7$. For both models, the sequential design continues up to the time horizon. Figure 3 compares the designs by computing the posterior expected TPs (gray bars) and the number of genes with limma P-values < 0.05 after the Benjamini and Yekutieli (2001) adjustment (black bars). At the time horizon both quantities increase over 2-fold compared with the recommended fixed sample size. The differences between prior and posterior expected TPs show how sequential designs adapt to the observed data to correct prior mis-specifications.

5.4 *Reverse-phase protein arrays*

We design a follow-up study for the RPPA dataset `dataIII` that is included in the R package `RPPanalyzer` (Mannsperger *and others*, 2010). The data contain expression for 75 proteins and 35 stage 2 and 25 stage 3 samples. Both models, the NN and GaGa models provide a reasonable fit (supplementary material available at *Biostatistics* online, Section 3). The fit under the NN model is slightly better. We find $\hat{\pi} = 0.13$ under the NN model, and $\hat{\pi} = 0.10$ under the GaGa model; that is, the estimated number of DE proteins is 9.75 and 7.5, respectively. Although we expect several DE proteins, at a posterior expected FDR < 0.05 the NN model calls one DE protein, and the GaGa model makes no DE calls. For comparison, only one protein has limma BY-adjusted P-values below 0.05.
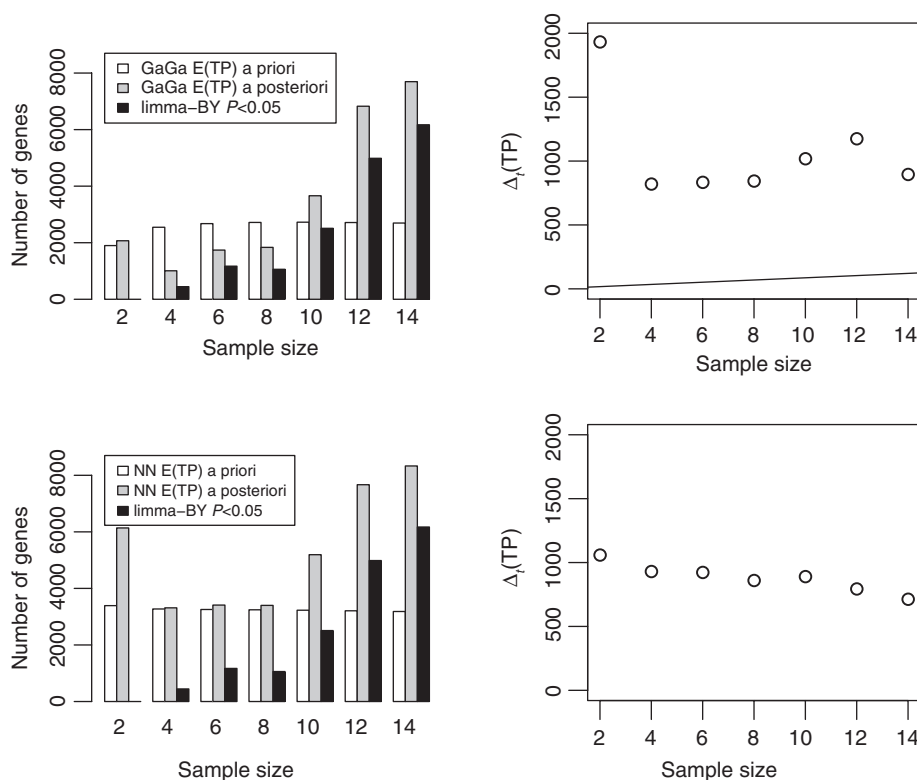
Fig. 3. Sequential Analysis of Campo Dell'Orto's data based on GaGa (top panels) and NN models (bottom panels). Left panels: expected number of TP (*a priori* and *a posteriori*) vs. sample size. Black bars indicate the number of genes with Benjanimi–Yekutieli-adjusted limma P-values < 0.05. Right panels: $\Delta_t(TP)$ vs. sample size and optimal sequential boundary. $\Delta_t(TP)$ being above the boundary for all t indicates the experimentation continuing up to the maximum sample size.

We consider adding batches of 50 samples per group, up to a maximum of $T = 4$. We set the sampling cost to $c = 1$, reflecting that RPPA samples are relatively cheap. The study focuses on 75 carefully chosen proteins. Figure 4 shows simulated $\Delta_t(TP)$ trajectories and the optimal boundaries. Inference under the GaGa model estimates fewer TPs. Otherwise, inference is fairly similar across models and the optimal boundaries are remarkably close. The average sample size is $t_S^* = 1.37$ under the NN model, and $t_S^* = 1.32$ under the GaGa model. The expected number of TPs at t_S^* is 9.46 (NN) and 6.11 (GaGa); that is, according to both models, most DE proteins should be detected by adding 1–2 batches, i.e. 50–100 samples per group. These results help assess the potential benefits of extending the experiment.

6. DISCUSSION

We proposed a sequential strategy for massive multiple hypotheses testing. An important advantage lies in the generality of the proposed design. We discussed three RNA-seq, one microarray, and one RPPA experiments. Sequential designs are robust with respect inaccurate prior guesses and provide substantial advantages over fixed sample designs.

The proposal is formulated in a decision-theoretic framework and emphasizes interpretability. We monitor the one-step ahead expected increment in utility and stop the experiment when it falls below a boundary.

Fig. 4. Simulated $\Delta_t(\text{TP})$ and optimal boundaries for $c = 1$ in RPPA data using GaGa (left) and NN (right) models.

The approach includes fixed sample size and myopic designs as special cases. We use terminal decisions that control the posterior expected FDR. While inconsistent with a strict decision-theoretic setup, where all decisions are taken to maximize the expectation of a single utility, we feel that our choice offers a pragmatic compromise.

The method allows stopping when only one or two samples are available, which requires making strong parametric assumptions. For instance, in Figure 3, the posterior expected TPs and $\Delta_t(\text{TP})$ based on two samples per group differ widely between the GaGa and NN models. Nevertheless, both models correctly indicate to continue and show good agreement for subsequent samples. Whenever possible, we recommend using a minimum burn-in (e.g. $\geqslant 3$ samples) before starting sequential stopping. When not feasible, we recommend assessing the goodness-of-fit carefully and updating the forward simulation when more data are available.

We focused on group comparison experiments, but the framework can serve as the basis for other HTSs. Interesting extensions include classification, clustering, or network discovery studies. These would require adjusting the utility function and possibly the probability model, e.g. to capture strong dependencies between outcomes.

Sequential designs are most appealing in moderate to large studies, where technical limitations require gathering data in batches. They should also prove valuable when samples are costly to obtain or there are ethical considerations, e.g. in human studies. Overall, they help save valuable resources and guarantee that sufficient data are collected to answer the scientific questions.

SOFTWARE

An implementation of the proposed approach was added to the Bioconductor package gaga.

SUPPLEMENTARY MATERIAL

Supplementary material is available at http://biostatistics.oxfordjournals.org.

ACKNOWLEDGMENTS

References

Armstrong, S. A., Staunton, J. E., Silverman, L. B., Pieters, R., Boer, M. L., Minden, M. D., Sallan, E. S., Lander, E. S., Golub, T. R. and Korsmeyer, S. J. (2002). MLL translocations specify a distinct gene expression profile that distinguishes a unique leukemia. *Nature Genetics* **30**, 41–47.

Benjamini, Y. and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics* **29**, 1165–1188.

Berry, D. A. and Fristedt, B. (1985). *Bandit Problems*. London: Chapman & Hall.

Berry, D. A., Müller, P., Grieve, A. P., Smith, M., Parke, T., Blazed, R., Mitchard, N. and Krams, M. (2001). *Adaptive Bayesian Designs for Dose-Ranging Drug Trials*, Volume V. New York: Springer.

Brockwell, A. E. and Kadane, J. B. (2003). A gridding method for Bayesian sequential decision problems. *Journal of Computational and Graphical Statistics* **12**, 566–584.

Campo Dell'Orto, M., Zangrando, A., Trentin, L., Li, R., Liu, W. M., te Kronnie, G., Basso, G. and Kohlmann, A. (2007). New data on robustness of gene expression signatures in leukemia: comparison of three distinct total rna preparation procedures. *BMC Genomics* **8**, 188.

Carlin, B., Kadane, J. and Gelfand, A. (1998). Approaches for optimal sequential decision analysis in clinical trials. *Biometrics* **54**, 964–975.

DeGroot, M. (2004). *Optimal Statistical Decisions*. Hoboken, NJ: Wiley-Interscience.

Dobbin, K. and Simon, R. M. (2007). Sample size planning for developing classifiers using high-dimensional DNA microarray data. *Biostatistics* **8**, 101–117.

Durrieu, G. and Briollais, L. (2009). Sequential design for microarray experiments. *Journal of the American Statistical Association* **104**, 650–660.

Ferreira, J. A. and Zwinderman, A. (2006). Approximate power and sample size calculations with the Benjamini–Hochberg method. *The International Journal of Biostatistics* **2**, 8.

Gibbons, R. D., Bhaumik, D. K., Cox, D. R., Grayson, D. R., Davis, J. M. and Sharma, R. P. (2005). Sequential prediction bounds for identifying differentially expressed genes in replicated microarray experiments. *Journal of Statistical Planning and Inference* **129**, 19–37.

Kenneth M. R., Thomas, L. and Szpiro, A. A. (2008). Trading bias for precision: decision theory for intervals and sets. *Technical Report*, Working Paper 336, UW Biostatistics, http://www.bepress.com/uwbiostat/paper336.

Lee, M. L. and Whitmore, G. (2004). Power and sample size for microarray studies. *Statistics in Medicine* **11**, 3543–3570.

Mannsperger, H. A., Gade, S., Henjes, F., Beissbarth, T. and Korf, U. (2010). RPPanalyzer: analysis of reverse-phase protein array data. *Bioinformatics* **26**, 2202–2203.

Müller, P., Berry, D., Grieve, A., Smith, M. and Krams, M. (2006). Simulation-based sequential Bayesian design. *Journal of Statistical Planning and Inference* **137**, 3140–3150.

Müller, P., Parmigiani, G., Robert, C. and Rousseau, J. (2004). Optimal sample size for multiple testing: the case of gene expression microarrays. *Journal of the American Statistical Association* **99**, 990–1001.

Pan, W., Lin, J. and Le, C. T. (2002). How many replicates of arrays are required to detect gene expression changes in microarray experiments? *Genome Biology* **3**, research0022.1–0022.10.

D. ROSSELL AND P. MÜLLER

ROSSELL, D. (2009). GaGa: a simple and flexible hierarchical model for differential expression analysis. *Annals of Applied Statistics* **3**, 1035–1051.

ROSSELL, D., MÜLLER, P. AND ROSNER, G. (2006). Screening designs for drug development. *Biostatistics* **8**, 595–608.

RUPPERT, D., NETTLETON, D. AND HWANG, J. T. G. (2007). Exploring the information in p-values for the analysis and planning of multiple-test experiments. *Biometrics* **63**, 483–495.

SMYTH, G. K. (2004). Linear models and empirical Bayes methods for assessing differential expression in microarray experiments. *Statistical Applications in Genetics and Molecular Biology* **3**, Number 1, Article 3.

TIBSHIRANI, R. (2006). A simple method for assessing sample sizes in microarray experiments. *BMC Bioinformatics* **7**, 106.

YANG, Y. H. AND SPEED, T. (2002). Design issues for cDNA microarray experiments. *Nature Reviews Genetics* **3**, 579–588.

YUAN, M. AND KENDZIORSKI, C. (2006). A unified approach for simultaneous gene clustering and differential expression identification. *Biometrics* **62**, 1089–1098.

ZIEN, A., FLUCK, J., ZIMMER, R. AND LENGAUER, T. (2003). Microarrays: how many do you need? *Journal of Computational Biology* **10**, 653–667.