

Survival Analysis of Longitudinal Microarrays

Natasa Rajcic, Dianne M. Finkelstein, David A. Schoenfeld

and the Inflammation and Host Response to Injury Research Program
Investigators¹
Massachusetts General Hospital, Biostatistics Unit, and Department of
Biostatistics, Harvard School of Public Health, Boston, MA

May 2006

Abstract

Motivation: The development of methods for linking gene expressions to various clinical and phenotypic characteristics is an active area of genomic research. Scientists hope that such analysis may, for example, describe relationships between gene function and clinical events such as death or recovery. Methods are available for relating gene expression to measurements that are categorized or continuous, but there is less work in relating expressions to an observed event time such as time to death, response, or relapse. When gene expressions are measured over time, there are methods for differentiating temporal patterns. However, no methods have yet been proposed for the survival analysis of longitudinally collected microarrays.

Results: We describe an approach for the survival analysis of longitudinal gene expression data. We construct a measure of association between the time to an event and gene expressions collected over time. The issue of high dimensionality and dependence when assessing statistical significance is addressed using permutations and control of the false discovery rate. Our proposed method is illustrated on a data set from a multi-center research study of inflammation and response to injury that aims to uncover the biological reasons why patients can have dramatically different outcomes after suffering a traumatic injury (www.gluegrant.org).

Contact: dfinkelstein@partners.org

1 Introduction

Scientists are turning to the microarray technology for insights into the mechanisms of the human body that were previously poorly understood. We think of genes as units of heredity as they record the genetic makeup of organisms. Though it is believed that a large number of genes remain inactive for most of our lives, there are those genes for which the activity can be associated with various physiological or environmental effects. In simple terms, a gene is

¹Contributing members of the Inflammation and the Host Response to Injury investigators are listed in Acknowledgments.

considered to be activated, or *expressed*, if its coded information is converted into proteins which are the main instigators of functions and processes in our bodies. An interesting problem in the analysis of the human genome is to relate changes in gene activity to clinical or phenotypic information. For example, scientists have been able to relate gene expression to the clinical implications of different types of cancer [Alizadeh *et al.*, 2000, Golub *et al.*, 1999, van't Veer *et al.*, 2002, van de Vijver *et al.*, 2002].

The nature of the data generated from a microarray experiment poses specific challenges to the statistical analysis. In a typical experiment, data from a relatively small number of subjects is available on thousands, even tens of thousands of genes, which makes many of the classical statistical procedures unapplicable. This is because the standard statistical methods are developed with the classical data type in mind, where the number of explanatory variables does not exceed the number of subjects on which data is collected. Time-to-an-event data poses additional challenges due to the presence of censoring.

We propose a method to study relationships between repeatedly collected gene expressions and time to an event of interest. Our problem is motivated by the data from Inflammation and Host Response to Injury research project (also referred to as the Glue Grant, [Glue Grant, 2006]). This multi-center and multi-disciplinary collaboration aims at better understanding of processes involved in the immune system's response to injury, as well as uncover the biological reasons why seemingly alike patients can have dramatically different outcomes after suffering a traumatic injury or burns. Doctors hope that identifying the genetic factors will help predict the course of recovery of severely injured patients.

Numerous methods have been proposed and developed for relating gene microarrays to either continuous or categorical measurements such as comparison of treatment groups or the level of a known biomarker. [Ring & Ross, 2002] offer a comprehensive review of methods that use microarrays for tumor classification. In comparison, fewer methods have been suggested for the analysis of gene expressions in relation to time to an event (e.g., death, response, or relapse), or for detecting differences in genes over time [Luan & Li, 2002, Yeung *et al.*, 2003, Storey *et al.*, 2005]. Methods for the use of longitudinal microarray in either describing or predicting survival outcomes are currently unavailable.

Methods for relating gene activity to the occurrence of an event have been proposed when microarrays are collected at a single point in time (e.g., baseline). One approach is to first use an unsupervised classification method, e.g. hierarchical clustering, to generate two or more groups of patient samples [Rosenwald *et al.*, 2002, Makretsov *et al.*, 2004]. The survival distributions within such generated clusters are then compared using the logrank test and displayed by Kaplan-Meier curves. A second, related approach is to first cluster genes based on their expressions across different patient samples, and then use cluster averages of the gene expressions as explanatory variables in a Cox proportional hazard regression model [Li & Luan, 2003]. However, both of these approaches do not capture the marginal relationship between gene expressions and time to an event [Sorlie *et al.*, 2001, Jung *et al.*, 2005]. One may end up with gene classes that do not represent any meaningful grouping in terms of the survival, or the results

may vary due to a particular clustering algorithm used.

A number of published approaches apply partial least squares (PLS) method to generate linear combinations of gene expressions as predictors in the proportional hazards model [Nguyen & Rocke, 2002]. PLS is a method related to principal components analysis (PCA), but while PCA creates combinations that maximize the explained variability among predictors only, PLS aims at maximizing correlations between predictors and the response variable. [Bair & Tibshirani, 2004] first calculate the Cox score for each gene (a statistic based on a proportional hazards partial likelihood) in order to select a subset of genes, then employ the PLS method to a reduced set of genes to arrive at the best model. [Park *et al.*, 2002] first reformulate the survival outcomes problem into a generalized linear (Poisson) regression, then apply the PLS algorithm to derive a parsimonious model. In a related approach, [Li & Gui, 2004] and [Gui & Li, 2004] reduce the dimensionality of the microarray predictor space by either partial or penalized Cox regression. While in all of these approaches the high-dimensionality of the microarray is reduced, direct interpretation of the fitted parameters in terms of the individual genes is not possible. In contrast, we are interested in developing a method to be used as a first step in identifying individual genes for further investigation.

When thousands of genes are measured in a single experiment, a single question of interest can be formulated as simultaneous testing of numerous individual hypotheses. Testing many statistical hypotheses at once increases the possibility of a Type I error, as a significant result may occur purely by chance, regardless of the nature's true state. The control of the false discovery rate (FDR) has become a widely used method of error control in the analysis of gene microarrays [Storey & Tibshirani, 2003]. The control of FDR involves an estimate of the proportion of falsely positive genes among all genes found *positive* (i.e., exhibit differential expression in different samples or states under investigation). [Westfall & Young, 1989] promoted the use of permutations to allow for dependencies among test statistics. In this paper, we propose a permutation-based method that is related to the method of [Storey & Tibshirani, 2003] and to the popular SAM method [Tusher *et al.*, 2001].

A gene-specific test statistic is defined in Section 2.2. A multiple testing algorithm that controls the number of false positive findings is described in Section 2.4. The results of a series of simulations are presented in Section 3, while the analysis on the data from the study of inflammation and response to trauma is presented in Section 4. All programs for the analysis presented in this paper were done using R statistical package ([R, 2005]) and can be obtained from the first author (N.R.) or from <http://hedwig.mgh.harvard.edu/biostatistics/software.php>.

2 Methods

2.1 Notation

For each of the n patients enrolled in a clinical study, we record the elapsed time from the beginning of the study to the occurrence of the event of interest (e.g., death or recovery). Since we are interested in examining whether there is an association between the time to an event and changes in gene expression, we also observe patients' microarray over time. A microarray here is a collection of p expression values on p different genes. For a given gene, let X_{ij} denote an expression for patient i at time t_j . Time to occurrence of the event is recorded using two indicator variables, δ_{ij} and Y_{ij} , where $\delta_{ij} = 1$ if subject i has experienced the event at time t_j , and $\delta_{ij} = 0$ otherwise. Similarly, $Y_{ij} = 1$ if i is still *at risk* at time t_j , meaning that the patient is under observation and has not experienced the event by time t_j . A set of subjects remaining at risk at time t_j is of size $Y_j = \sum_i Y_{ij}$. Let $n_j = \sum_i \delta_{ij}$ denote a total number of subjects with the event at time t_j . For patients for which the event does not occur for the duration of the study, or who for other reasons have discontinued the study followup, we say that their event time is censored.

2.2 Test statistic

To test the association between thousands of longitudinally collected gene expressions and time to an event of interest, we want to use a test statistic that is not only intuitive but simple to calculate. This is primarily because we intend to use a permutation-based testing procedure which is computationally intensive. The approach we take is to first calculate one test statistic per each gene, then determine the significance of each association using permutations. We begin by examining a general class of nonparametric tests for survival data, formulated by [Jones & Crowley, 1989]. We assume that for each subject i at risk at time j , it is possible to define a (quantitative) value Z_{ij} that represents subject's covariate measurement, and denote by \bar{Z}_j the average value of Z for subjects at risk at time t_j . It is also assumed that the increasing covariate values correspond to either increasing or decreasing chances of event occurrence. Let m denote the total number of *observed* events, then the test statistic can be written as:

$$T(\omega, Z) = \sum_{j=1}^m \omega_j \sum_{i=1}^n \delta_{ij} [Z_{ij} - \bar{Z}_j]. \quad (1)$$

Note that ω_j are optional weights chosen to emphasize either early or late events. Since it is assumed each subject can experience the event only once (i.e., absorbing state), it follows that m also represents the total number of subjects with events. The score statistic based on the partial likelihood from a Cox regression model [Cox, 1972] is a member of this general class. [Jones & Crowley, 1989] investigate various choices for weights and labels (ω_j, Z_{ij}) , where using $(1, X_{ij})$

results in the following test statistic:

$$T = \sum_{j=1}^m \sum_{i=1}^n \delta_{ij} [X_{ij} - \bar{X}_j]. \quad (2)$$

Here, \bar{X}_j is the average gene expression for subjects at risk at time t_j , $\bar{X}_j = (1/Y_j) \sum_i X_{ij}$. This test statistic captures the difference between the observed covariate values for subjects that had an event at a given time-point, and the average covariate value for subjects still at risk an instant before the event occurred. The differences are then summed up over all observed event times.

In our approach, however, we cannot implement (2) without further modification. This is because the structure of our problem presents several challenges that need to be addressed. The outer summation in both (1) and (2) is over unique observed event times, which results in summands involving current gene expression values at the time of an observed event. Ideally, if we had observed expression values for all subjects currently at risk at time of an observed event, the statistic would be easily and correctly calculated. In clinical trials, however, data are often collected according to some schedule of study visits, so that the data on time-varying covariates may not be available at time of an event for all subjects at risk at that time. For example, in the Glue study, gene expression is obtained on 7 scheduled visits over a period of 28 days. A simple approach to deal with intermittent covariate data is the "last observation carried forward" approach (LOCF), where the most recent available observation is used in place of the missing data. If the event had occurred at some time t_* between the two scheduled visits t_{j-1} and t_j , so that X_{i*} is not available at the time of the event, microarray collected at time t_{j-1} would be used in place of X_{i*} . While the last observation carried forward approach would be simple to implement, it has been traditionally heavily criticized as it produces biased results. We therefore explore a different approach to handling intermittent microarray data.

2.3 Semi-parametric test of association

Another way to deal with the intermittently available expression data is to model unknown values using measurements available up to that time. Taking into account our limited knowledge about the longitudinal behavior of microarrays, we search for ways to model the gene expression over time without assuming strict distributional properties. Following the idea of a conditional score [Stefanski & Carroll, 1987], [Tsiatis & Davidian, 2001] model longitudinal and event data using a proportional hazards regression model for the relationship between the event hazard at time t_j , λ_{ij} , and the longitudinal covariate, X_{ij} : $\lambda_{ij} = \lambda_0 \exp(\beta X_{ij})$, where λ_0 is an underlying event hazard, and β is a parameter associating the hazard and the covariate. The longitudinal process is modelled via the random effects model, $X_{ij} = \alpha_{0i} + \alpha_{1i} t_j + \epsilon_{ij}$, where $\alpha_i = (\alpha_{0i}, \alpha_{1i})$ are subject-specific random intercept and slopes, and ϵ_{ij} is a mean-zero normally distributed random variable representing the measurement error. By applying the random effects to model individual subject trajectories, we also address the

issue of between-subject gene expression variability which can be substantial (Cheng Li, personal communication).

The fitted value of X_{ij} from the random effects model that uses all available data up to and including time t_j is denoted by $\hat{X}_{ij} = \hat{\alpha}_{0i} + \hat{\alpha}_{1i}t_j$. In the literature, random intercepts and slopes, $\alpha_i = (\alpha_{0i}, \alpha_{1i})$, are usually assumed to be normally distributed. However, if we know the sufficient statistic for α_i , we can use it to avoid the normality assumption. Namely, the sufficient statistic for α_i , conditional on subject i being at risk at time t_j , i.e., $Y_{ij} = 1$, is

$$S_{ij} = \beta\sigma_j^2\delta_{ij} + \hat{X}_{ij}.$$

Consequently, the joint likelihood of the events δ_{ij} and the model estimate \hat{X}_{ij} can be factored into two parts, one of which does not involve a random variable α_i , and an other that does not involve information on the event:

$$L(\delta_{ij}, \hat{X}_{ij}|\alpha_i) = L(\delta_{ij}|S_{ij}) \times L(\hat{X}_{ij}|\alpha_i).$$

The conditional likelihood $L(\delta_{ij}|S_{ij})$ does not depend on the random effect α_i , and arguably contains all the relevant information about the parameter of interest, β . The error term, σ_j^2 , accounts for the uncertainty when using \hat{X}_{ij} as an estimate of the unknown covariate value X_{ij} . We describe the estimation of σ_j^2 below. [Tsiatis & Davidian, 2001] use this conditional likelihood to suggest the following estimating equations for parameter β :

$$\sum_{j=1}^m \sum_{i=1}^n S_{ij} [\delta_{ij} - \exp(\beta S_{ij} - \frac{\beta^2 \sigma_j^2}{2})] Y_{ij} = 0. \quad (3)$$

Rather than estimating β , we construct a score test statistic based on the estimating equations. We then use the test statistic to test the null hypothesis $H_0 : \beta = 0$ vs. $H_1 : \beta \neq 0$. Similar to the generalized score test statistic described by [Rotnitzky & Jewell, 1990], we define the test statistics based on (3) as $T = U(0)/\sqrt{\text{Var}[U(0)]}$, where

$$U(0) = \sum_{j=1}^m \sum_{i=1}^n \delta_{ij} [\hat{X}_{ij} - \frac{\sum_i \hat{X}_{ij} Y_{ij}}{\sum_i Y_{ij}}]. \quad (4)$$

The factor in (3) that is quadratic in β disappears in the score statistic which is the first derivative of (3) evaluated at $\beta = 0$. Also note that this score test statistic has a familiar form to that in (2). The difference between the two statistics is that (4) estimates the unknown value of the covariate at the observed event time using the available covariate history. To obtain an estimate of variance, we calculated the second derivative of the score function and evaluated it for $\beta = 0$,

$$\hat{\text{Var}}[U(0)] = \sum_{j=1}^m \frac{n_j}{Y_j - 1} [\hat{\sigma}_j^2(Y_j - n_j) + (Y_j - 1)\hat{V}_j]. \quad (5)$$

The estimate $\hat{\sigma}_j^2$ is a product of two quantities. First is the estimate of the unknown variability in measuring X_{ij} at t_{ij} , which is estimated by the pooled estimate of the residual sums of squares over all subjects. The second quantity is the estimate of variance of the predicted value \hat{X}_{ij} using available values up to and including time t_j . Specifically, variance of the predicted value \hat{X}_{ij} at time t_j is $1/m_{i,j} + (t_j - \bar{t}_{i,j})^2/SS_{i,j}$, where $m_{i,j}$ is the number of available observations for subject i up to time t_j , and $SS_{i,j}$ is the corresponding sum of squared differences from the mean, using values up to an including time t_j . As before, Y_j is the total number of subjects at risk at time t_j , and n_j is the number of events at time t_j . $\hat{V}_j = \frac{1}{Y_j} [\sum \hat{X}_i^2 - Y_j \hat{X}^2]$ is the sample variance of the ordinary least square estimates, \hat{X}_{ij} , among subjects at risk at t_j . The resulting score test closely resembles the score test from the Cox's proportional hazards model for non-time-varying covariates or for completely known time-varying covariate histories. Our proposed test can also be viewed as a test of form $H_0 : \lambda(t|X) = \lambda(t)$, for all X , where $\lambda(t)$ is a hazard function.

2.4 Calculation of the significance levels

The statistic in (4) is best suited for examining a single covariate, whereas we need to test thousands of genes (covariates) to determine which gene's changes over time are associated with a clinical event. Moreover, when testing genes simultaneously, we want to take into account the possible (but unknown) dependencies among genes. To determine the significance of the test statistics, we propose a permutation-based testing procedure. Our approach removes the need for explicitly specifying the joint distribution of thousands of gene-specific test statistics because our permutation algorithm simulates the correct asymptotic multivariate distribution. This point is explained further at the end of this section.

At each observed event time, we permute the event indicators among subjects at risk at that time. In other words, the number of subjects with events is kept fixed at each event time, but their event indicators randomly exchanged among those currently at risk. In what follows, we outline the algorithm as if all tests statistics, both those calculated on the original data and those calculated on the perturbed data, are positive. We do this for the ease of presentation, otherwise only the notation would become more involved. Let T_1, \dots, T_p be test statistics calculated on each of the p -genes in the original data, then proceed as follows:

1. At each observed event time t_j , permute indicators of events among those subjects still at risk at that time; This is equivalent to choosing n_j elements out of Y_j , at time t_j ;
2. Using such perturbed data, calculate new test statistics, labelled T^* , one per each gene: T_k^* , $k = 1 \dots p$;
3. Compare each original T_k with all permutation-based T^* and call the

number of *false positives* the number among T^* that are greater than T_k ,

$$FP(T_k) = \sum(T^* > T_k), \quad k = 1, \dots, p; \quad (6)$$

4. Repeat steps 1-3 many times, say, hundred times. For each gene k , $k \in \{1, \dots, p\}$, this produces a sequence of hundred numbers. Denote by $\overline{FP}(T_k)$ the mean value of such sequence for gene k ;
5. For each gene, the estimated proportion of false positives is the ratio of $\overline{FP}(T_k)$ over all statistics called significant for that gene. The number of remaining T considered significant for T_k is the number that exceed in absolute value: $\sum(T > T_k)$. Thus, the estimated proportion of false positives for T_k is:

$$\frac{\overline{FP}(T_k)}{\sum(T > T_k)} \quad k = 1, \dots, p \quad (7)$$

If a test statistic has an estimated proportion of false positives below a desired, pre-specified level, say 10%, then the hypothesis is rejected and the test statistic is declared statistically significant.

Note that the permutations are performed in order to estimate the number of false positive, $FP(T_k)$, per each calculated test statistic T_k (Step 3). This quantity is also used to estimate the numerator of the proportion of false positives (Step 5). If we knew the multivariate distribution of the test statistics, $FP(T_k)$ could be calculated by numerical integration which, in the case of thousands of test statistics, is only possible by simulation. However, [Davidov & Zelen, 1998] show that the asymptotic multivariate distribution of the score test statistics can be approximated using the urn sampling scheme, which is what we *de facto* do. Thus, the algorithm we present uses the permutation-based simulation of the asymptotic multivariate joint distribution that has the correct multivariate covariance structure. As a consequence, our estimate of false positive (numerator in FDR) appropriately accounts for the unknown dependencies among the score test statistics.

We calculate the permutation distribution by permuting the event indicators among subjects at risk at each time. Alternatively, the true permutation distribution for the test statistics would ostensibly be found by permuting the event times among the patients. The problem with this approach is that no samples were collected after the event for each patient, and if they were, the gene expression values after the event may have been affected by the event which would preclude their use. One way to fill-in such missing data, would be to define a distance metric in order to select among subjects with non-missing data those that are 'close' or 'similar' to the subject with the missing observation. The algorithm will then proceed as follows: a) cluster subjects according to their microarray at time t_{k-1} , b) note the cluster membership of the subject with the missing t_k array, and c) impute the missing array by calculating some sample measure (e.g. mean array) using the remaining members of the cluster and their expressions at time t_k . The testing algorithm can then continue with the Step 2 above.

3 Simulations

We performed a series of simulations to assess validity and performance of our proposed method. The following describes an algorithm to generate longitudinal expressions along with survival outcomes that emulate the data-generating mechanisms presented by the actual problem.

1. Obtain an estimate of a variance-covariance matrix, $\hat{\Sigma}$, of the random effects by fitting a random effects model to a selection of genes from the actual data;
2. Sample from a bivariate normal distribution with zero mean and variance-covariance matrix obtained in Step 1 to get a set of random effects (intercepts and slopes): $(a_{0i}, a_{1i}) \sim N_2(\mathbf{0}, \hat{\Sigma})$, for $i = 1, \dots, n$;
3. Generate individual gene trajectories for each subject and for all genes, using the generated random effects: $X_{ij} = a_{0i} + a_{1i}t_j + \epsilon_{ij}$, where $\epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$;
4. Choose a value of the association parameter b . Assuming an exponential hazard function $\lambda(t) = \lambda_0 e^{bX(t)}$, the parameter b captures the strength of association between time to an event and the time-dependent covariate trajectory $X(t)$. Using the exponential form of the hazard function $\lambda(t)$, and an estimate of an underlying event hazard, λ_0 , which uses the number of observed events, we derive an inverse of the cumulative hazard distribution:

$$\Lambda^{-1}(u) = \frac{1}{ba_1} \log\left[\frac{-\log(u) ba_1}{\hat{\lambda}_0 e^{ba_0}} + 1\right]; \quad (8)$$

5. By knowing the form of the inverse cumulative hazard function, we can use the Probability Integral Transformation to sample survival times, T ; We first generate replicates of the uniformly distributed random variable $U \sim Unif(0, 1)$, then generate survival times as $T = \Lambda^{-1}[-\log(U)]$.

In the final step of the algorithm, generated survival times that exceed 28 days are considered censored. The produced survival times are 'linked' to the trajectories generated in Step 3 through the random effects α . Namely, since the same random effects generated in Steps 1-2 are used in Step 5 to generate survival times T , subjects with comparable random effects get assigned similar event times (e.g., early or late).

Using the above algorithm, we generated 600 samples of data. Each sample consists of 100 subjects with 500 longitudinal gene expressions over 7 time-points. Fifty out of 500 genes were set to be significantly associated with the time to an event. Testing was done for three choices of the association parameter b , as well as two values for the measurement error, σ_ϵ^2 . Within each simulation,

Table 1: Simulation results
 $n = 100$ subjects; $p = 500$ genes; 600 replications

median (IQR)	b= 2.5	b= 1.5	b= 0 prop.positive
$\sigma_\epsilon^2 = 0.10$			
# positive	56 (53, 58)	55 (53, 57)	0.173
prop. false +	0.107 (0.056, 0.136)	0.090 (0.057, 0.137)	–
$\sigma_\epsilon^2 = 0.20$			
# positive	55 (52.75, 57)	53 (48, 55)	0.141
prop. false +	0.092 (0.056, 0.137)	0.090 (0.063, 0.125)	–

the false positive ratio of 10% was used as a cutoff value for determining significance. Simulations were executed using R statistical software and results presented in Table 1.

To understand these results, let us examine the case when the association parameter b is set to 2.5, the measurement error of individual genes is $\sigma_\epsilon^2 = 0.20$, and 50 out of a total of 500 simulated genes have trajectories associated with the time to an event (i.e., 10% of genes are significant). A median number of genes found positive over 600 simulations is 55 genes. The median false positive proportion over 600 simulations is 0.092, with an interquartile range of (0.056, 0.137). Similar results are found for the remainder of the cells. When $b = 0$, we report the proportion of (falsely) positive genes, as all genes are expected to be non-significant under $H_0 : \beta = 0$. If the association parameter b is set to zero, any significant genes should be found purely by chance, and we would expect the total number of significant genes to be zero.

Inspection of the simulation results shows that our method performs reasonably well. The proportion of false positive findings remains close to the pre-specified 10% mark for both choices of the association parameter b and the two levels of measurement error. Also note that the proportion of false positives remains similar across the two columns. A change of the pre-specified association parameter b does not dramatically change the proportion of false positive. However, the assumed level of the measurement error seems to influence the estimated number of false positive findings. Estimates of false positive proportions are higher for $\sigma_\epsilon^2 = 0.10$.

4 Results of Trauma Data Analysis

We apply our method to the data from Inflammation and Host Response to Injury research project (the Glue study). This collaborative program examines the biological reasons why patients can have dramatically different clinical outcomes after experiencing a traumatic injury. Among many scientific questions posed by the Glue investigators is whether we can identify genes whose tempo-

ral changes relate to the time until a specific clinical event. It is reasonable to assume that genes exhibiting greater variation are more likely to be associated with the time to an event. Patients in the Glue study are followed for 28 days from the time they experience a serious traumatic injury. This is an example of a right-censored data, which we assume for the developments of our method. Genomic data collected on days 0, 1, 4, 7, 14, 21, and 28 are generated using commercially available oligonucleotide array technology [Affymetrix Inc., 2001]. Each microarray includes expressions on 54,674 probe sets (which we will call 'genes' for the purposes of our analysis). Gene expressions were extracted from oligonucleotide probes by employing a *PM-only* analysis of [Li & Wong, 2001] and normalized across arrays to achieve comparable levels, using the 'Invariant Set' method in the dChip software ([Li & Wong, 2003]). Finally, the gene expression values were log-transformed prior to any calculations.

To reduce the overwhelming dimensionality of a microarray, we first excluded those genes labelled 'Absent' over all arrays by the Affymetrix software. We then performed a simple filtering of genes and included only those genes whose estimated coefficient of variation (CV) exceeds a certain threshold. While more complex filtering can be used, one can also proceed without filtering at all. Our choice of threshold is somewhat ad-hoc as we aimed at having a couple of thousands of genes to work with, instead of over fifty four thousand, in this hypothesis-generating approach. This brought the number of genes to under four thousand ($p = 3,914$). Data on 56 subjects with complete entries were included in the analysis. For the purposes of survival analysis, we define the event of interest as "respiratory recovery" which occurs when a patient no longer depends on a machine to breathe. The response is thus defined as the time from injury (and entry into the ICU) to getting off the ventilator. Of 56 patients, 52 experienced recovery of respiratory function prior to the end of 28-day follow-up. Four patients remained on the ventilator by day 28 and thus had recovery time censored at 28 days.

The test statistic that measures the association between each gene and the time to an event is calculated for each gene separately. We performed described permutation-based testing procedure to determine significance of each test statistic. Of 3,914 investigated genes, 154 were identified as statistically significant when the proportion of false positive findings is set at the 10% level. As a comparison, a total of 694 genes were identified as significantly associated with the time to a recovery when their test statistics are compared to the 10 th -percentiles of the normal distribution.

A sample gene ontology for a selection of genes for which change in expression over time is associated with the time to respiratory recovery is given in Table 2.

Results for a sample of four significant genes are presented in Figure 1, with plots of individual patient trajectories over time. The four panels are identified by accession numbers. In order to further illustrate our results, patients are distinguished by whether they had a 'late' respiratory recovery (those occurring after day 16). Solid lines represent a patients that either experienced a recovery after day 16 or their times were censored at day 28. The two plots on the right-hand side are for genes negatively related to the time to a recovery. A decrease in gene expression on these two genes

Table 2: Gene ontology for a subset of significant genes

Accession #	Function (association direction)
NM_002668	has a role in chemotactic processes via CCR1 (-)
NM_004994	linked to increased invasiveness of cancer cells (-)
NM_001629	associated with myocardial infarction and stroke (-)
NM_002800	fatty acid biosynthesis and oxidation (+)

is related to a shorter time to recovery. The opposite is true for the other two plots where an increase in gene expression relates to a shorter time to recovery.

The tens of thousands of gene expressions measured repeatedly over time on tens or hundreds of patients can create a considerable computational difficulty due to the enormity of the resulting data sets. To further help reduce the time needed for lengthy computations, we collaborated with an application specialist at the Massachusetts General Hospital Biostatistics Unit in order to obtain, develop, test, and employ a parallel computing system [Lazar & Schoenfeld, 2004]. The 30+ node computer cluster helped us greatly reduce the time needed for lengthy computations.

5 Discussion

We provide a method for the survival analysis of longitudinally collected microarrays. We address the issues of intermittently collected microarray data, unknown longitudinal behavior of a single gene expression, as well as joint distribution of thousands of genes. Our testing procedure is similar to the approach proposed by [Storey & Tibshirani, 2003], when the estimated proportion of null hypotheses $\hat{\pi}_0$ is set to 1. Our procedure is also equivalent to the popular SAM procedure when symmetric thresholds are used [Tusher *et al.*, 2001]. In our analysis, we were not interested in explicitly monitoring the Type II error. However, as one of the reviewers pointed out, an appropriate Type II error control is difficult in this setting and should be carefully considered.

More complex longitudinal models can easily be incorporated into our approach. A minimum set of assumptions regarding the functional relationship between longitudinal gene expression and timing of the events will depend on an individual biological problem at hand. For example, a natural extension would be to implement the approach of [Song *et al.*, 2002], which requires only the assumption that the random effects have a smooth density. Another modification, which may be relevant in some applications, is to devise a multiple imputation procedure for the unknown covariate values. While this will certainly add to the overall computational complexity, it would be interesting to explore whether it can be incorporated so to take advantage of the computations already in place and the high-dimensionality of the data. Finally, in order to make the proposed test statistics more robust to potential outliers, the actual values of gene expressions may have be replaced by ranks or some function of the ranks.

In the Glue study, patients were closely monitored at all times for a period of 28 days. The study subjects either experience an event or their time is censored at the end of the study follow-up. This is an example of Type I censoring where there is no possibility of missing data due to a dropout. However, in a typical clinical trial where study participants are followed for a longer period of time, it is likely that the censoring

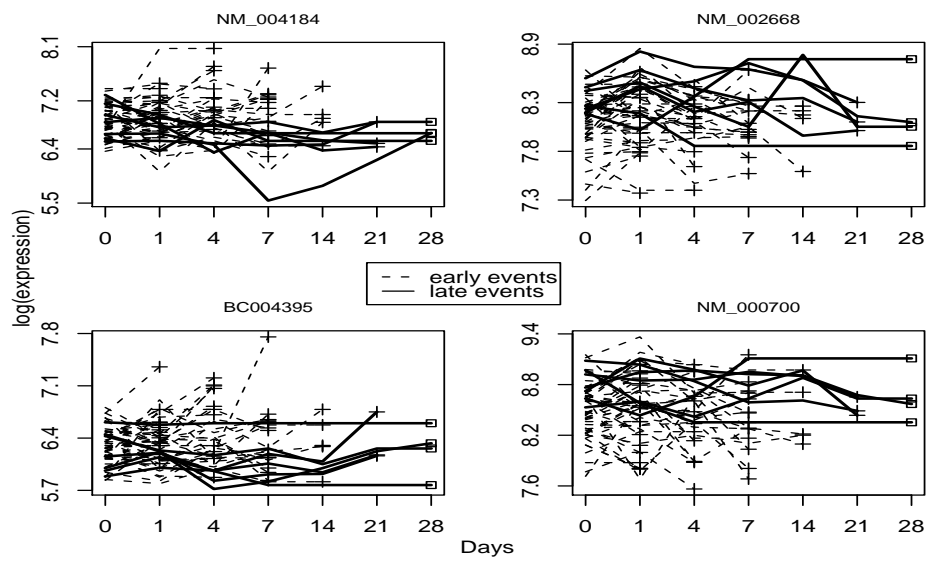


Figure 1: Plot of patients trajectories for four genes found significant. Solid lines represent patients with recovery events that occurred after day 16, or patients censored at day 28. The test statistics corresponding to the two plots to the right are negative, indicating an inverse relationship between gene expressions and the time to a recovery. Conversely, the test statistics corresponding to the two plots to the left are positive, an increase in gene expression is associated with a shorter time to recovery.

due to early dropout would be an issue. It is straightforward to accommodate this type of censoring in our approach.

The continual advancement of the microarray technology will ultimately result in many large studies routinely including longitudinal genomic observations as part of the study follow-up. The longitudinal microarray and event time data will thus become more common. Also, other applications of high-dimensional data such as proteomics and metabolomics data will arise. Therefore, further development of efficient methodologies to handle these high-dimensional event time data sets will be needed.

Acknowledgement

We wish to thank Dr. Cheng Li for valuable insights on the subject. We also extend our thanks to Drs. John Storey and John Quackenbush for their helpful suggestions and review. The application specialist, Peter Lazar, engineered and maintained the parallel system in R used in our analysis. Finally, we thank the four reviewers whose careful review and constructive criticism greatly improved the quality of the manuscript.

The following is a list of participating investigators of the Inflammation and the Host Response to Injury research project: Henry V. Baker, Ronald V. Maier, Paul E. Bankey, John A. Mannick, Timothy R. Billiar, Philip H. Mason, Bernard H. Brownstein, Grace P. McDonald-Smith, Steve E. Calvano, Bruce A. McKinley, David G. Camp II, Carol Miller-Graziano, Celeste Campbell-Finnerty, Michael N. Mindrinos, George Casella, Joseph P. Minei, Irshad H. Chaudry, Lyle L. Moldawer, Mashkoor Choudhry, Ernest E. Moore, J. Perren Cobb, Frederick A. Moore, Ronald W. Davis, Avery B. Nathens, Asit De, Grant E. O'Keefe, Constance Elson, Laurence G. Rahme, Bradley Freeman, Daniel G. Remick, Richard L. Gamelli, David A. Schoenfeld, Nicole S. Gibran, Michael B. Shapiro, Douglas L. Hayden, Martin Schwacha, Brian G. Harbrecht, Geoffrey M. Silver, David N. Herndon, Richard D. Smith, Jureta W. Horton, John Storey, William Hubbard, Ronald G. Tompkins, John Lee Hunt, Mehmet Toner, Jeffrey Johnson, H. Shaw Warren, Matthew B. Klein, Michael A. West, James A. Lederer, Wenzhong Xiao, Tanya Logvinenko, Stephen F. Lowry.

References

- [Affymetrix Inc., 2001] Affymetrix, Inc. 2001. Microarray Suite 5.0. Santa Clara, CA.
- [Alizadeh *et al.*, 2000] Alizadeh AA, Eisen MB, Davis RE, et al. 2000. Distinct types of diffuse large Bcell lymphoma identified by gene expression profiling. *Nature* 403, 503-11.
- [Bair & Tibshirani, 2004] Bair, E., Tibshirani R. 2004. Semi-supervised methods for predicting patient survival from gene expression papers. *PLoS Biology* 2, 5011-22.
- [Cox, 1972] Cox DR. 1972. Regression models and life tables. *J R Stat Soc B34*, 187-20.
- [Davidov & Zelen, 1998] Davidov O, Zelen M. 1998. Urn sampling and the proportional hazard model. *Lifetime Data Analysis* 4, 309-27.
- [Golub *et al.*, 1999] Golub TR, Slonim DK, Tamayo P, et al. 1999. Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science* 286, 531-37.
- [Gui & Li, 2004] Gui L, Li H. 2005. Penalized Cox regression analysis in the high-dimensional and low-sample size settings, with applications to microarray gene expression data. *Bioinformatics*, 21(13):3001-08.
- [Glue Grant, 2006] Inflammation and Host Response to Injury: a multi-disciplinary research project. 2006. National Institute of General Medical Sciences (NIGMS) <http://www.gluegrant.org>.
- [Jones & Crowley, 1989] Jones M, Crowley J. 1989. A general class of nonparametric tests for survival analysis. *Biometrics* 45, 15770.

- [Jung *et al.*, 2005] Jung SH, Owzar K, George SL. 2005. A multiple testing procedure to associate gene expression levels with survival. *Stat Med*, 24(20):3077-88.
- [Lazar & Schoenfeld, 2004] Lazar P & Schoenfeld D. 2004. Self-contained parallel system for R. MGH Biostatistics, Boston, MA. <http://cran.r-project.org/src/contrib/Descriptions/biopara.html>
- [Li & Luan, 2003] Li H, Luan Y. Kernel Cox regression models for linking gene expression profiles to censored survival data. 2003. University of California, Davis, Proc Pac Symp Biocomp, 65-76.
- [Li & Gui, 2004] Li H, Gui L. 2004. Partial Cox regression analysis for high-dimensional microarray gene expression data. *Bioinformatics*, 20(Suppl. 1).
- [Li & Wong, 2001] Li C., Wong WH. 2001. Model-based analysis of oligonucleotide arrays: Expression index computation and outlier detection. *Proc Natl Acad Sci*, 98(1), 31-36.
- [Li & Wong, 2003] Li C, Wong WH. 2003. DNA-Chip Analyzer (dChip). In The analysis of gene expression data: methods and software. Edited by Parmigiani G, Garrett ES, Irizarry R, and Zeger SL. Springer.
- [Luan & Li, 2002] Luan Y, Li H. 2002. Clustering of time-course gene expression data using a mixed-effects model with B-splines. *Bioinformatics*, 19(4):474-82.
- [Makretsov *et al.*, 2004] Makretsov NA, Huntsman DG, Nielsen TO, et al. 2004. Hierarchical clustering analysis of tissue microarray immunostaining data identifies prognostically significant groups of breast carcinoma. *Clin Cancer Res*, 10: 6143-51.
- [Nguyen & Rocke, 2002] Nguyen D, Rocke D. 2002. Partial Least Squares Proportional Hazard Regression for Application to DNA Microarray Survival Data. *Bioinformatics* 18(12): 1625-32.
- [Park *et al.*, 2002] Park P, Tian L, Kohane I. Linking Gene Expression Data with Patient Survival Times Using Partial Least Squares. 2002. *Bioinformatics*, 18(1):S120-27.
- [R, 2005] R Development Core Team. R: A Language and Environment for Statistical Computing. 2005. R Foundation for Statistical Computing, Vienna, Austria, <http://www.R-project.org>.
- [Ring & Ross, 2002] Ring BZ, Ross DT. 2002. Microarrays and molecular markers for tumor classification. *Genome Biology*, 3:comment 2005.1-6.
- [Rosenwald *et al.*, 2002] Rosenwald A, Wright G, Chan WC, *et al.* 2002. The use of molecular profiling to predict survival after chemotherapy for diffuse large-B-cell lymphoma. *N Engl J Med*, 346:1937-47.
- [Rotnitzky & Jewell, 1990] Rotnitzky A, Jewell N. 1990. Hypothesis testing of regression parameters in semiparametric generalized linear models for cluster correlated data. *Biometrika*, 77:485-97.
- [Song *et al.*, 2002] Song X, Tsiatis AA, Davidian M. 2002. A Semiparametric likelihood approach to joint modeling of longitudinal and time-to-event data. *Biometrics*, 58:742-53.
- [Sorlie *et al.*, 2001] Sorlie T, Perou CM, Tibshirani R *et al.* 2001. Gene expression patterns of breast carcinomas distinguish tumor subclasses with clinical implications. *Proc Natl Acad Sci*, 98(19):10869-74.
- [Stefanski & Carroll, 1987] Stefanski LA, Carroll RJ. 1987. Conditional scores and optimal scores in generalized linear measurement error models. *Biometrika* 74, 703-16.
- [Storey & Tibshirani, 2003] Storey JD, Tibshirani R. 2003. Statistical significance for genomewide studies. *Proc Natl Acad Sci*, 100(16):9440-45.
- [Storey *et al.*, 2005] Storey JD, Xiao W, Leek JT, Tompkins RG, Davis RW. 2005 Significance analysis of time course microarray experiments. *Proc Natl Acad Sci*, 102(36):12837-42.
- [Tsiatis & Davidian, 2001] Tsiatis AA, Davidian M. 2001. A semiparametric estimator for the proportional hazards model with longitudinal covariates measured with error. *Biometrika*, 88(2):447-458.

- [Tusher *et al.*, 2001] Tusher V, Tibshirani R, Chow G. 2001. Significance analysis of microarrays applied to the ionizing radiation response. *Proc Natl Acad Sci*, 98:5116-21.
- [van't Veer *et al.*, 2002] van't Veer LJ, Dai H, van de Vijver MJ, et al. 2002. Gene expression profiling predicts clinical outcome of breast cancer. *Nature*, 415:530-36.
- [van de Vijver *et al.*, 2002] van de Vijver MJ, He YD, van 't Veer LJ, et al. A gene-expression signature as a predictor of survival in breast cancer. *N Engl J Med*, 347:1999-2009.
- [Westfall & Young, 1989] Westfall P, Young S. 1989. p Value Adjustments for Multiple Tests in Multivariate Binomial Models. *J Am Stat Assoc*, 84(407):780-86.
- [Yeung *et al.*, 2003] Yeung KY, Medvedovic M, Bungarner RE. 2003. Clustering gene-expression data with repeated measurements. *Genome Biology*, 4(5):R34.