Gaussian process regression for survival time prediction with genome-wide gene expression

AARON J. MOLSTAD

Biostatistics Program, Public Health Sciences Division, Fred Hutchinson Cancer Research Center, 1100 Fairview Ave. N., Seattle, WA, USA

LI HSU

Biostatistics Program, Public Health Sciences Division, Fred Hutchinson Cancer Research Center, 1100 Fairview Ave. N., Seattle, WA, USA and Department of Biostatistics, University of Washington, 1705 NE Pacific St, Seattle, WA 98195, USA

WEI SUN*

Biostatistics Program, Public Health Sciences Division, Fred Hutchinson Cancer Research Center, 1100 Fairview Ave. N., Seattle, WA, USA and Department of Biostatistics, University of Washington, 1705 NE Pacific St, Seattle, WA 98195, USA and Department of Biostatistics, University of North Carolina at Chapel Hill, 135 Dauer Dr, Chapel Hill, NC 27599, USA

wsun@fredhutch.org

SUMMARY

Predicting the survival time of a cancer patient based on his/her genome-wide gene expression remains a challenging problem. For certain types of cancer, the effects of gene expression on survival are both weak and abundant, so identifying non-zero effects with reasonable accuracy is difficult. As an alternative to methods that use variable selection, we propose a Gaussian process accelerated failure time model to predict survival time using genome-wide or pathway-wide gene expression data. Using a Monte Carlo expectation–maximization algorithm, we jointly impute censored log-survival time and estimate model parameters. We demonstrate the performance of our method and its advantage over existing methods in both simulations and real data analysis. The real data that we analyze were collected from 513 patients with kidney renal clear cell carcinoma and include survival time, demographic/clinical variables, and expression of more than 20 000 genes. In addition to the right-censored survival time, our method can also accommodate left-censored or interval-censored outcomes; and it provides a natural way to combine multiple types of high-dimensional -omics data. An R package implementing our method is available in the [Supplementary material](#) available at *Biostatistics* online.

Keywords: Gaussian process accelerated failure time model; Gene expression; Multiple kernel learning; Survival time.

*To whom correspondence should be addressed.

Fig. 1. Density plots of the marginal p -values for each of the 20 483 genes, gene expression effect sizes, and concordance for the 20 483 marginal models in plots (a), (b), and (c) for Model I, and plots (d), (e), and (f) for Model II. Dark densities are from the observed data, whereas white densities are for permuted gene expression data with the patient labels of gene expression data permuted. Dark dotted lines in (c) and (f) denote the concordance measurement of survival time prediction under the baseline model that includes all the covariates in Model I or II, other than gene expression.

1. INTRODUCTION

Predicting the survival time of a cancer patient based on his/her genome-wide gene expression is a well-studied, yet unresolved problem. The effects of gene expression could be both weak and abundant which, when combined with often high censoring rates, make feature selection for survival time association very challenging. On the other hand, genome-wide gene expression data can be highly informative for prognosis. For example, [Zhu and others \(2017\)](#) demonstrate that two patients with similar genome-wide gene expression data may have similar survival time.

Our method development is motivated by a data set of more than 500 cancer patients with kidney renal cell carcinoma, which is part of The Cancer Genome Atlas (TCGA) project. We seek to predict survival time based on genome-wide gene expression and demographical/clinical variables. To demonstrate that the associations between gene expression and survival time are abundant and weak in this data set, we first report results of gene-by-gene marginal association testing. For each gene, we fit two Cox proportional hazards models. In Model I, we include only the expression of this gene as a predictor and the sequencing plate ID as a confounder. In Model II, we include the expression of this gene, sequencing plate ID, and three demographical/clinical covariates: age, gender, and tumor stage. We evaluate the marginal effect of each gene using marginal p -value (Figures 1(a) and (d)) and effect size (Figures 1(b) and (e)), and compare the distributions of these association statistics using observed data (dark densities) versus permuted data (light densities). The contrast of these two distributions clearly demonstrates the abundant and weak effects of gene expression on survival time.

Assuming genes with p -value larger than 0.5 are not associated with survival time, we can calculate the expected number of genes associated with survival time by $p - 2s$, where s is the number of genes

with p -value > 0.5 and $p = 20483$ is the total number of genes. This number is 13 052 and 10 316 for Models I and II, respectively. Similarly, with a false discovery rate of 0.05, the number of genes that are significantly associated with survival time are 8550 and 4002 for Models I and II, respectively. The large number of genes associated with survival time is biologically plausible given that kidney renal clear cell carcinoma is characterized by oncogenic metabolism and epigenetic reprogramming, both of which may affect the expression of many genes ([Cancer Genome Atlas Research Network, 2013](#)).

In Figures 1(c) and (f), we display the survival time prediction concordances (C-index) for each of the 20483 genes. The dark dashed vertical lines denote the concordance of the baseline model, i.e., the model excluding gene expression but including sequencing plate ID (Model I) as well as clinical covariates (Model II). For Model I, including a single gene can improve concordance by as much as 6%. Comparatively, the improvement in concordance for Model II is smaller, with a single gene improving concordance no more than 2%. The concordance improvements indicate that gene expression can improve the prediction of survival time, although few, if any, genes appear to have strong effects. Together, these results suggest that screening or variable selection may be difficult or ineffective because of potentially weak and abundant effects.

The majority of methods for survival time prediction address censoring using partial likelihood methods, which use event orderings rather than the times at which they occur. Consequently, when survival time can be predicted with reasonable accuracy, partial likelihood methods may miss useful information in the censoring times. Alternatively, some methods use a two-step approach to first impute censored survival times (e.g., mean, median, or multiple imputation), and then fit a predictive model using the imputed survival times ([Datta and others, 2007](#); [Wu and others, 2008](#)).

Some other methods iteratively impute the censored survival times and fit a predictive model, for example, using survival trees ([Zhu and Kosorok, 2012](#)) or an ensemble model ([Deng and others, 2016](#)). These methods were not designed for high-dimensional -omic data. For example, in their real data examples, the sample size (n) and the number of covariates (p) are $n = 686$ and $p = 8$ for [Zhu and Kosorok \(2012\)](#), and $n = 2070$ and $p = 256$ for [Deng and others \(2016\)](#). In contrast, we propose a new method that iteratively imputes the censored survival times and fits a kernel-based predictive model using multiple kernels. Our real data analysis with $n = 513$ and $p = 20428$ has much higher dimensionality than the earlier methods. [Zhu and others \(2017\)](#) employed a kernel-based method for survival time prediction using gene expression; however, they only used one kernel derived from gene expression and did not seek to impute the censored survival times.

In our proposed method, we do not attempt to identify a subset of genes associated with survival time. Instead, we use genome-wide gene expression to model the covariance of the log-survival time under a Gaussian process accelerated failure time (AFT) model. Inspired by multiple kernel learning ([Gönen and Alpaydm, 2011](#)), we allow the covariance to be a linear combination of M user-specified candidate kernels, allowing for greater flexibility to capture abundant and weak effects than previous approaches. A major challenge for survival time prediction is censoring. To mitigate this challenge, we develop a novel and efficient Monte Carlo expectation-maximization (MC-EM) algorithm which jointly imputes censored log-survival times and estimates model parameters. The imputed survival times are then used in our subsequent prediction.

The remainder of this article is organized as follows: in Section 2, we describe our proposed model and discuss its relation to existing methods; in Section 3, we describe how to compute our estimator; in Section 4, we perform simulation studies to demonstrate our method's prediction accuracy under a range of models; in Section 5, we analyze the TCGA data set which motivated our study; and in Section 6, we discuss limitations and extensions of our method.

2. GAUSSIAN PROCESS ACCELERATED FAILURE TIME MODEL

Let S_i denote the time-to-failure (survival time) for the i th patient with $i = 1, \dots, n$ patients in the study. Let $T = (\log S_1, \dots, \log S_n)' \in \mathbb{R}^n$. Let $x_i \in \mathbb{R}^p$ and $z_i \in \mathbb{R}^{q+1}$ denote the measured genome-wide gene expression and the measured clinical variables for the i th patient, respectively. To allow for an intercept, assume that the first entry of z_i is equal to one for $i = 1, \dots, n$. Let $Z = (z_1, \dots, z_n)' \in \mathbb{R}^{n \times (q+1)}$ and $X = (x_1, \dots, x_n)' \in \mathbb{R}^{n \times p}$. For the n patients in the study, we assume that survival time follows the *Gaussian process accelerated failure time model*:

$$T = Z\beta + G + \epsilon, \quad G \sim N_n\{0, K(X, \sigma^2)\}, \quad \epsilon \sim N_n\{0, \sigma_\epsilon^2 I_n\}, \quad (2.1)$$

where G and ϵ are independent; $\sigma^2 \in \mathbb{R}_+^M$, $\sigma_\epsilon^2 > 0$, and $\beta \in \mathbb{R}^{q+1}$ are unknown model parameters, $\mathbb{R}_+$ denotes non-negative real numbers, and M is the number of kernels. We will sometimes use the more compact notation: $\text{Cov}(G + \epsilon) \equiv \tilde{K}(X, \tilde{\sigma}^2) = K(X, \sigma^2) + \sigma_\epsilon^2 I_n$, where $\tilde{\sigma}^2 = (\sigma^2', \sigma_\epsilon^2)' \in \mathbb{R}_+^{M+1}$.

The function $K : \mathbb{R}^{n \times p} \times \mathbb{R}_+^M \rightarrow \mathbb{S}_+^n$ is a covariance function with (i, j) th entry

$$[K(X, \sigma^2)]_{ij} = \sum_{s=1}^M \sigma_s^2 k_s(x_i, x_j), \quad (i, j) \in \{1, \dots, n\} \times \{1, \dots, n\},$$

where $\mathbb{S}_+^n$ denotes the set of $n \times n$ symmetric and positive semidefinite matrices, and $k_s : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ is a positive semidefinite kernel function for $s = 1, \dots, M$. A positive semidefinite kernel function ensures that the matrix $k_s(X, X) : \mathbb{R}^{n \times p} \times \mathbb{R}^{n \times p} \rightarrow \mathbb{S}_+^n$, whose (i, j) th entry is $k_s(x_i, x_j)$, is positive semidefinite for all $X \in \mathbb{R}^{n \times p}$. The function $k_s(x_i, x_j)$ quantifies the similarity between x_i and x_j , e.g., a radial basis kernel function is $k_s(x_i, x_j) = \exp(-\|x_i - x_j\|^2)$.

The Gaussian process AFT model in (2.1) generalizes the log-normal AFT model of [Klein and others \(1999\)](#), which for clustered subjects, assumed that $\text{Cov}(T_i, T_j) = \phi$ for all (i, j) such that i and j belong to the same cluster and $i \neq j$. A familiar special case occurs when $K(X, \sigma^2) = \sigma^2 X'X$, in which case (2.1) would be equivalent to a linear mixed model with normal priors on the regression coefficients for gene expression. Model 2.1 is also connected to the kernel machine model with possibly non-linear function of G that belongs to the function space generated by a kernel function $K(\cdot, \cdot)$ ([Liu and others, 2007](#)). Gaussian processes have also been used for survival analysis under the Cox proportional hazards model ([Banerjee and others, 2003](#); [Fernández and others, 2016](#); [Zhu and others, 2017](#)). Intuitively, (2.1) assumes that if two patients have similar genome-wide gene expression, as defined by K , then their mean-adjusted log-survival times will be similar. Out-of-sample prediction based on (2.1) is also known as *kriging*, a method for prediction through linear interpolation in geo-spatial statistics. In geo-spatial applications, the function K is used to quantify the similarities of 2D coordinates, whereas in our application, K quantifies similarities in an ultra-high dimensional, genome-wide space. Recently, kriging was applied in the genomic literature as a means for predicting a phenotypic trait using multiple types of -omics data ([Wheeler and others, 2014](#)).

Fitting (2.1) is non-trivial when one observes a censored realization of T , as is often the case in survival analysis. Specifically, suppose there exists a realization $T = (t_1, \dots, t_n)' \in \mathbb{R}^n$, which cannot be observed. Instead one observes the pairs $(y_1, \delta_1), \dots, (y_n, \delta_n)$ where $y_i = \min(t_i, d_i)$, d_i is the log censoring time for the i th subject, $\delta_i = 1(y_i = t_i)$ for $i = 1, \dots, n$, and $1(\cdot)$ is an indicator function. In this article, we treat the censored survival times as missing. This allows us to develop an algorithm that simultaneously imputes the latent survival times conditional on the observed survival times and model parameters; and estimates model parameters $\beta, \sigma^2, \sigma_\epsilon^2$. Although, we focus on the case of right-censored outcomes, our methodology also naturally accommodates left or interval censoring.

For the remainder of the article, without loss of generality, suppose that $\delta_i = 0$ for $i = 1, \dots, n_c$, $\delta_i = 1$ for $i = n_c + 1, \dots, n$, and let $n_o = n - n_c$. Hence, we can partition $Y = (y_1, \dots, y_n)'$ into $Y_c \in \mathbb{R}^{n_c}$ and $Y_o \in \mathbb{R}^{n_o}$ so that $Y = (Y'_c, Y'_o)' \in \mathbb{R}^n$. We similarly partition T into $(T'_c, T'_o)'$ (where T_c is not observed and $T_o = Y_o$); Z into $Z_c \in \mathbb{R}^{n_c \times (q+1)}$ and $Z_o \in \mathbb{R}^{n_o \times (q+1)}$; and $\tilde{K}(X, \tilde{\sigma}^2)$ into submatrices $\tilde{K}_{co}(X, \tilde{\sigma}^2) \in \mathbb{R}^{n_c \times n_o}$, $\tilde{K}_{oo}(X, \tilde{\sigma}^2) \in \mathbb{R}^{n_o \times n_o}$, and $\tilde{K}_{cc}(X, \tilde{\sigma}^2) \in \mathbb{R}^{n_c \times n_c}$. For ease of display, we will sometimes omit the $(X, \tilde{\sigma}^2)$ dependence on $\tilde{K}(X, \tilde{\sigma}^2)$ and its submatrices. Let $\mathcal{H} = \mathbb{R}^{q+1} \times \mathbb{R}_+^M \times \mathbb{R}_+$ denote the space of the unknown parameters $\theta = (\beta', \sigma^{2'}, \sigma_\epsilon^2)'$. Finally, let W be the collection of data that we condition on: $W = \{Z, X, Y, \delta\}$.

3. MAXIMUM LIKELIHOOD ESTIMATION

3.1. Overview

To fit the Gaussian process AFT model in (2.1), we use a MC-EM algorithm. We provide an overview this algorithm in Section 3.2 and describe the key sub-algorithm in Section 3.3. We implement our method in an R package `SurvGPR`, which is available at <https://github.com/ajmolstad/SurvGPR>.

3.2. Monte Carlo expectation–maximization algorithm

Throughout this section, let the superscript (r) denote the r th iterate of the MC-EM algorithm, and let s_r denote the r th iterate's Monte Carlo sample size. The $(r + 1)$ th iterate of the standard expectation–maximization (EM) algorithm is computed in two steps: the E-step computes

$$Q(\theta \mid \theta^{(r)}) = \mathbb{E}[\log f_T(T_o, T_c; \theta, W) \mid \theta^{(r)}, W], \quad (3.2)$$

where $\log f_T$ is the log-likelihood of T ; and the M-step computes

$$\theta^{(r+1)} = \arg \max_{\theta \in \mathcal{H}} Q(\theta \mid \theta^{(r)}). \quad (3.3)$$

When (3.3) cannot be obtained, an alternative is to compute $\theta^{(r+1)}$ such that

$$\theta^{(r+1)} \in \{\theta \in \mathcal{H} : Q(\theta \mid \theta^{(r)}) \geq Q(\theta^{(r)} \mid \theta^{(r)})\}, \quad (3.4)$$

which yields the generalized EM algorithm (Wu C.F.J, 1983).

Unfortunately, when log-survival times are censored, there may not exist an analytic expression for the right-hand side of (3.2) under (2.1). In particular, ignoring constants,

$$Q(\theta \mid \theta^{(r)}) \propto -\mathbb{E}[\log \det\{\tilde{K}(X, \tilde{\sigma}^2)\} + (T - Z\beta)' \{\tilde{K}(X, \tilde{\sigma}^2)\}^{-1} (T - Z\beta) \mid \theta^{(r)}, W], \quad (3.5)$$

so that computing $Q(\theta \mid \theta^{(r)})$ requires evaluating

$$(i) \mathbb{E}[T_c \mid \theta^{(r)}, W], \quad (ii) \mathbb{E}\left[T'_c (\tilde{K}_{cc} - \tilde{K}_{co} \tilde{K}_{oo}^{-1} \tilde{K}'_{co})^{-1} T_c \mid \theta^{(r)}, W\right].$$

Computing (i) and (ii) is non-trivial because

$$T_c \mid \theta^{(r)}, W \sim N_{n_c}^{[Y_c, \infty)} \left\{ Z_c \beta^{(r)} + \tilde{K}_{co} \tilde{K}_{oo}^{-1} (T_o - Z_o \beta^{(r)}), \tilde{K}_{cc} - \tilde{K}_{co} \tilde{K}_{oo}^{-1} \tilde{K}'_{co} \right\}, \quad (3.6)$$

where the notation $N_{n_c}^{[Y_c, \infty)}$ denotes the n_c -dimensional truncated multivariate normal distribution with non-zero probability mass on the hyper-rectangle $[Y_c, \infty) = [y_1, \infty) \times \cdots \times [y_{n_c}, \infty)$. Although (i) can be computed numerically, the distribution of $(\tilde{K}_{cc} - \tilde{K}_{co}\tilde{K}_{oo}^{-1}\tilde{K}'_{co})^{-1/2}T_c \mid (\theta^{(r)}, W)$ is not truncated multivariate normal unless $(\tilde{K}_c - \tilde{K}_{co}\tilde{K}_{oo}^{-1}\tilde{K}'_{co}) = I_{n_c}$ (Horrace, 2005), so (ii) is intractable. We approximate (3.2) by drawing s_r samples from (3.6) (Wei and Tanner, 1990).

There are multiple software packages available to simulate from (3.6). In our implementation, we use the Gibbs sampler implemented in the `tmvtnorm` package in R (Wilhelm and Manjunath, 2015). Let $\tilde{T}_c^{(r)} = (T_{c,1}^{(r)}, \dots, T_{c,s_r}^{(r)})' \in \mathbb{R}^{s_r \times n_c}$ be the matrix of samples from (3.6). Given $\tilde{T}_c^{(r)}$, the $(r+1)$ th iterate of our MC-EM algorithm is

$$\theta^{(r+1)} = \arg \max_{\theta \in \mathcal{H}} \left\{ s_r^{-1} \sum_{j=1}^{s_r} \log f_T(T_o, T_{c,j}^{(r)}; \theta, W) \right\}. \quad (3.7)$$

We propose an algorithm to compute (3.7) in Section 3.3. To improve the efficiency of our MC-EM algorithm, we use the ascent-based variation proposed by Caffo and others (2005). We state our complete ascent-based MC-EM algorithm in Algorithm 1.

Algorithm 1: Initialize $\theta^{(1)} = (\boldsymbol{\beta}^{(1)'}, \boldsymbol{\sigma}^{2(1)'}, \sigma_\epsilon^{2(1)'})'$. Set $r = 1$ and $s_1 = 500$.

(1) Simulate $\tilde{T}_c^{(r)}$, s_r samples from

$$N_{n_c}^{[Y_c, \infty)} \left\{ Z_c \boldsymbol{\beta}^{(r)} + \tilde{K}'_{co} \tilde{K}_{oo}^{-1} (T_o - Z_o \boldsymbol{\beta}^{(r)}), \tilde{K}_{cc} - \tilde{K}_{co} \tilde{K}_{oo}^{-1} \tilde{K}_{oc} \right\}.$$

(2) Compute $\bar{\theta} \leftarrow \arg \max_{\theta \in \mathcal{H}} \left\{ s_r^{-1} \sum_{j=1}^{s_r} \log f_T(T_o, T_{c,j}^{(r)}; \theta, W) \right\}$ (see Algorithm 2).

(3) Compute $\text{ASE}^{(r)}$, the standard error of $\left\{ \log f_T(T_o, T_{c,j}^{(r)}; \bar{\theta}, W) - \log f_T(T_o, T_{c,j}^{(r)}; \theta^{(r)}, W) \right\}_{j=1}^{s_r}$.

4a. If $s_r^{-1} \sum_{j=1}^{s_r} \left\{ \log f_T(T_o, T_{c,j}^{(r)}; \bar{\theta}, W) - \log f_T(T_o, T_{c,j}^{(r)}; \theta^{(r)}, W) \right\} > 1.96 \text{ASE}^{(r)}$

- Set $\theta^{(r+1)} \leftarrow \bar{\theta}$, $s_{r+1} = s_r$, $r \leftarrow r + 1$, and return to Step 1.

4b. Else

- If $s_r \geq 10^5$, terminate. Else, set $s_r \leftarrow 2s_r$ and return to Step 1, appending s_r new samples to the s_r from the previous iteration.

We terminate the algorithm based on a Monte Carlo sample size threshold in Step 4. When the algorithm has converged, the difference between $\theta^{(r)}$ and $\bar{\theta}$ will be negligible for sufficiently large s_r . In practice, we suggest practitioners track the values of the iterates to ensure that 10^5 is a sufficiently large threshold for their application.

Because we use a Gibbs sampler in Step 1, the simulated $T_{c,j}^{(r)}$ may be correlated. To decrease dependence while maintaining computational efficiency, we keep every tenth sample generated by the Gibbs sampler. To compute the standard error in Step 3 while accounting for the serial correlations due to the Gibbs sampler, we use the spectral variance method with a Tukey–Hanning window implemented in the R package `mcmcse` (Flegal and others, 2017).

3.3. Maximization algorithms

We now describe how to solve Step 2 of Algorithm 1. Throughout this section, treat r as fixed, let $\hat{T}_j = (T_{c,j}^{(r)'}, T_o')$ for $j = 1, \dots, s_r$, and let $\bar{T} = s_r^{-1} \sum_{j=1}^{s_r} \hat{T}_j$.

For all settings with $M \geq 1$, we use a variation of the blockwise coordinate descent algorithm proposed by [Zhou and others \(2018\)](#). The complete algorithm is stated in Algorithm 2, with explanations and details to follow.

Algorithm 2: Initialize $\theta^{(1)} = (\beta^{(1)'}, \sigma^{2(1)'}, \sigma_\epsilon^{2(1)'})'$ at their final iterates from the previous M-step. Set $b = 1$. Note this $\theta^{(1)}$ is different from the one in Algorithm 1 since it is indexed by b instead of r . We re-use the notation for notation simplicity.

- (1) Compute $\Omega \leftarrow \tilde{K}(X, \hat{\sigma}^{2(b)})^{-1}$
- (2) Compute $\beta^{(b+1)} \leftarrow (Z' \Omega Z)^{-1} Z' \Omega \bar{T}$
- (3) For $i = 1, \dots, M$, compute

$$\sigma_i^{2(b+1)} \leftarrow \frac{\sigma_i^{2(b)}}{\sqrt{s_r}} \left[\frac{\sum_{j=1}^{s_r} (\hat{T}_j - Z\beta^{(b+1)})' \Omega' k_i(X, X) \Omega (\hat{T}_j - Z\beta^{(b+1)})}{\text{tr} \{ \Omega k_i(X, X) \}} \right]^{1/2},$$

- (4) Compute

$$\sigma_\epsilon^{2(b+1)} \leftarrow \frac{\sigma_\epsilon^{2(b)}}{\sqrt{s_r}} \left[\frac{\sum_{j=1}^{s_r} (\hat{T}_j - Z\beta^{(b+1)})' \Omega' \Omega (\hat{T}_j - Z\beta^{(b+1)})}{\text{tr}(\Omega)} \right]^{1/2}.$$

- 5a. If $\{ \sum_{j=1}^{s_r} \log f_T(T_o, T_{c,j}^{(r)}; \theta^{(b+1)}, W) - \log f_T(T_o, T_{c,j}^{(r)}; \theta^{(b)}, W) \}$
 $\leq \epsilon | \sum_{j=1}^{s_r} f_T(T_o, T_{c,j}^{(r)}; \theta^{(1)}, W) |$
 - Terminate.
- 5b. Else
 - Set $b \leftarrow b + 1$ and return to Step 1.

The updates of $\sigma_i^{2(b+1)}$ and $\sigma_\epsilon^{2(b+1)}$ in Steps 3 and 4 are derived based on the minorize–maximize (MM) algorithm for fitting variance components models proposed by [Zhou and others \(2018\)](#). Briefly, given the initial values of the parameters or their estimates from the previous iteration, a minorizing function is created to approximate the objective function. The updates in Steps 3 and 4 are the arguments that maximize the minorizing function and thus, ensure that the objective function evaluated at $\theta^{(b+1)}$ is greater than or equal to the objective function evaluated at $\theta^{(b)}$. A complete derivation of Algorithm 2 is provided in the [Section 1 of the supplementary material](#) available at *Biostatistics* online.

In our implementation, we also adopt an extrapolation heuristic to accelerate convergence. We found that the iterates from Steps 3 and 4 of Algorithm 2 often followed monotonic paths to their local maximizers. Thus, after Step 4, we attempt to replace $\sigma^{2(b+1)}$ with an extrapolated value $\bar{\sigma}^{2(b+1)} = \sigma^{2(b+1)} + (b^{1/2} + 2)^{-1}(\sigma^{2(b+1)} - \sigma^{2(b)})$, and similarly for $\sigma_\epsilon^{2(b+1)}$. If the log-likelihood for the extrapolated values $\bar{\sigma}^{2(b+1)}$ and $\bar{\sigma}_\epsilon^{2(b+1)}$ is greater than the log-likelihood evaluated at the $\sigma^{2(b+1)}$ and $\sigma_\epsilon^{2(b+1)}$, we replace $\sigma^{2(b+1)}$ with $\bar{\sigma}^{2(b+1)}$ and $\sigma_\epsilon^{2(b+1)}$ with $\bar{\sigma}_\epsilon^{2(b+1)}$.

3.4. Implementation and practical considerations

Given the final iterates of the MC-EM algorithm, $\hat{\beta}, \hat{\sigma}^2, \hat{\sigma}_\epsilon^2$, and final imputed survival time, $\bar{T}$, we predict log-survival time for a new patient with covariates z_* and genome-wide gene expression x_* using the conditional expectation of the univariate normal distribution:

$$N \left\{ \hat{\beta}' z_* + K_*(x_*, X, \hat{\sigma}^2) \tilde{K}(X, \hat{\sigma}^2)^{-1} (\bar{T} - Z\hat{\beta}), \tilde{K}(x_*, \hat{\sigma}^2) - K_*(x_*, X, \hat{\sigma}^2)' \tilde{K}(X, \hat{\sigma}^2)^{-1} K_*(x_*, X, \hat{\sigma}^2) \right\},$$

where $K_*(x_*, X, \hat{\sigma}^2) \in \mathbb{R}^n$ with j th entry $[K_*(x_*, X, \hat{\sigma}^2)]_j = \sum_{s=1}^M \hat{\sigma}_s^2 k_s(x_*, x_j)$ for $j = 1, \dots, n$. We can also easily evaluate the estimated survival function, $\hat{S}$ at any time a since $P(T_* < a | z_*, x_*)$ is the cumulative distribution function of a univariate normal distribution.

In studies collecting gene expression or other types of -omics data, there are often technical confounders, e.g., the plate id for an RNA sample. To address confounding in genome-wide gene expression under (2.1), we propose to compute the kernel functions k_s using the residuals from the multivariate regression of gene expression on the measured technical confounders.

To obtain reasonable initial values for our MC-EM algorithm with right-censored survival times, we suggest first imputing the censored log-survival times using the inverse-probability-weighted mean-imputation method proposed by Datta (2005), or when additional covariates can be incorporated, the imputation method proposed by Grimes and others (2018).

In Section 2 of the supplementary material available at *Biostatistics* online, we present a sensitivity analysis which suggests that our MC-EM algorithm is reasonably robust to initial values and that the Monte Carlo error is negligible at convergence.

4. SIMULATION STUDIES

4.1. Data generating models

To create simulation scenarios similar to our motivating data example, we use the observed gene expression data and clinical covariates of the 513 patients in the TCGA KIRC (kidney renal clear cell carcinoma) data set, and we simulate survival times for these patients. Specifically, we use the observed tumor stage and age as clinical covariates, and use the observed expression of $p = 20483$ genes to generate survival times from four distinct models. More information about how we prepared the TCGA KIRC data set is given in Section 5.1. For 500 independent replications, we generate $n = 513$ survival times and split the data into a training and testing set of size 413 and 100, respectively. We then fit the model to the censored training data and record the metrics described in Section 4.3. The data generating models we consider are:

Model 1: Gaussian process AFT model. Log-survival times are generated as a realization of the Gaussian process AFT model:

$$T = Z\beta + \eta + \gamma,$$

where $\gamma \sim N_n\{0, 0.5I_n\}$ and $\eta \sim N_n\{0, K(X, \sigma^2)\}$ with $K(X, \sigma^2)$ defined below and $\beta = (6.1, -0.5, -1.2, -2.0, -1 \times 10^{-5})$ where the columns of Z correspond to the intercept, indicators for tumor stage II, III, and IV, and age in days.

Model 2: Normal-logistic AFT model. Log-survival times are generated as a realization of the normal-logistic AFT model,

$$T = Z\beta + \eta + \kappa,$$

where $\kappa = (\kappa_1, \dots, \kappa_n)'$ with each κ_i follows independent and identically distributed logistic distribution such that $E(\kappa_i) = 0$ and $\text{Var}(\kappa_i) = 0.5$. Note that logistic distribution has much heavier tails than normal distribution. As in Model 1, $\eta \sim N_n\{0, K(X, \sigma^2)\}$ with $K(X, \sigma^2)$ defined below; and Z and β are the same as in Model 1.

Model 3: Logistic-logistic AFT model. Log-survival times are generated as a realization of the logistic-logistic AFT model,

$$T = Z\beta + \omega + \kappa,$$

where κ is generated in the same manner as in Model 2. To generate $\omega \in \mathbb{R}^n$, we generate $v_1, \dots, v_n$, n independent copies of $V_i \sim \text{Logistic}$ where $E(V_i) = 0$ and $\text{Var}(V_i) = 1$ for $i = 1, \dots, n$. Then, we set $(\omega_1, \dots, \omega_n)' = (v_1, \dots, v_n)' \{K(X, \sigma^2)\}^{1/2}$ so that $E(\omega_i) = 0$ and $\text{Cov}(\omega_i, \omega_j) = [K(X, \sigma^2)]_{ij}$.

Model 4: Cox proportional hazards model. We generate survival times from the mixed-effects Cox proportional hazards model with Gompertz baseline hazard (Bender and others, 2005). Let $W = \tilde{Z}\beta + \eta$ where $\beta = (0.1, 0.3, 0.9, 9 \times 10^{-5})$ with columns of $\tilde{Z}$ corresponding to indicators for tumor stage II, III, and IV, and age in days, respectively; and let $\eta \sim N_n\{0, K(X, \sigma^2)\}$ with $K(X, \sigma^2)$ defined below. Then following Bender and others (2005), we generated survival times as a realization of

$$S_i = \frac{1}{\alpha} \log \left[1 - \frac{\alpha \log(u_i)}{\lambda \exp(W_i)} \right], \quad i = 1, \dots, n$$

where $u_1, \dots, u_n$ are n independent realizations of a $\text{Uniform}(0, 1)$ random variable; and $\alpha = \pi(1200\sqrt{6})^{-1}$ and $\lambda = \alpha \exp(-0.5772 - \alpha 1400)$ are chosen to mimic the survival distribution in the real data set with mean 1400 and standard deviation 1200.

For Models 1–4, the i th subject's censoring time is drawn from an exponential distribution with mean $\{Q_{\tau_i}(\{S_j\}_{j=1}^n)\}^{-1}$, where Q_c denotes the c th quantile and $\tau_i = 0.20, 0.50, 0.70$, or 0.80 for subjects with tumor stages I, II, III, or IV respectively. About 62% of survival times are censored on average in each of the four data generating models.

Models 2–4 illustrate our method's performance under three different types of misspecification. In Model 2, the error distribution is misspecified by our method, whereas in Model 3, both the genomic effect and the error distribution are misspecified by our method. Under Model 4, the log-linearity of survival time is violated.

For each of the four data generating models, we consider two variations of covariance function $K(X, \sigma^2)$:

Genome-wide kernel. We compute $K(X, \sigma^2)$ using a normalized radial basis function kernel based on genome-wide gene expression. Given $x_i \in \mathbb{R}^p$ for $i = 1, \dots, n$, we compute

$$[K(X, \sigma^2)]_{ij} = \sigma_G^2 \underline{k}(x_i, x_j) \equiv \sigma_G^2 \exp \left\{ \frac{-\|x_i - x_j\|^2}{\max_{l,m} (\|x_l - x_m\|^2)} \right\}. \quad (4.8)$$

In Models 1–3, we set $\sigma_G^2 = 3$, and in Model 4, we set $\sigma_G^2 = 4$. A histogram showing the off-diagonal entries of (4.8) is displayed in Figure 6(b) of the supplementary material available at *Biostatistics* online.

Pathway kernel. We compute $K(X, \sigma^2)$ as the sum of normalized radial basis function kernels as in (4.8) based on a small number of genes meant to emulate gene-pathways:

$$[K(X, \sigma^2)]_{ij} = \sum_{s=1}^6 \sigma_s^2 \underline{k}(D_s x_i, D_s x_j), \quad (i, j) \in \{1, \dots, n\} \times \{1, \dots, n\}, \quad (4.9)$$

where for $s = 1, \dots, 6$, $D_s \in \mathbb{R}^{p \times p}$ has 150, 150, 100, 100, 50, and 50 non-overlapping, randomly positioned ones on its diagonal and zeros in all other positions. Four of six σ_s^2 are randomly

assigned to be non-zero and their magnitudes are drawn independently from Uniform[0, 1] and normalized so that $\sum_{s=1}^6 \sigma_s^2 = 3$ in Models 1–3, and $\sum_{s=1}^6 \sigma_s^2 = 4$ in Model 4.

The genome-wide kernel data generating model illustrates the performance of the competing methods when effect sizes are small and abundant, i.e., genome-wide. The pathway-kernel data generating model is meant to compare our method to those which perform variable selection or marginal screening since there will be at most 500 genes that affect the survival time distribution.

4.2. Methods

In their review of methods for predicting survival time based on gene expression, [Van Wieringen and others \(2009\)](#) concluded that among the tree-based ensemble methods and regularized Cox proportional hazard models they compared, the L_2 -penalized Cox proportional hazards model performed as well as or better than the other methods. For this reason, we exclude tree-based ensemble methods from our comparisons, but include the L_1 and L_2 -penalized Cox proportional hazards model using genome-wide gene expression, tumor stage, and age as covariates. We do not penalize coefficients corresponding to tumor stage or age in either model and select tuning parameters using 10-fold cross-validation.

We also consider the regularized AFT models proposed by [Grimes and others \(2018\)](#) which perform simultaneous imputation and model fitting. Like the penalized Cox proportional hazards model, we use both L_1 and L_2 -penalized versions of the method of [Grimes and others \(2018\)](#).

Screening is often used when employing penalized models in ultra-high dimensional settings. Thus, we also consider L_1 and L_2 -penalized Cox models or AFT models using a pre-screened gene sets, which we call L_1^* and L_2^* . When the data generating model uses the genome-wide kernel, the pre-screened methods use only genes that are marginally associated with survival at the 0.10 significance level after adjusting for multiple testing. For the methods of [Grimes and others \(2018\)](#), screening is performed using imputed survival times. When the data generating model uses the pathway kernels, the screening method retains the 600 genes that are used to construct the pathway kernels, i.e., assuming an oracle screening method.

We use two variations of our Gaussian process AFT model. The versions denoted GPR : K and GPR : M correspond to the genome-wide kernel and pathway kernels, respectively. For GPR : K, we use (4.8) as the lone candidate kernel. For GPR : M, we use seven candidate kernels: the six pathway kernels from (4.9), and a genome-wide kernel which is similar to (4.8), except it uses all genes not used in any of the pathway kernels. When the data generating model uses the genome-wide kernel, the six pathway kernels are computed from 600 randomly selected genes with the same pathway sizes as the true pathway-kernels. When the data generating model uses the pathway-kernels, GPR : M includes the true pathway-kernels and genome-wide (excluding genes in the pathways) kernel as candidates.

To illustrate the benefit of our imputation procedure, which jointly imputes censored survival times and estimates model parameters, we also compare our method to two versions of the Gaussian process AFT model that first imputes survival times using the imputation procedure from [Grimes and others \(2018\)](#), which incorporates only the clinical covariates but not gene expression, and then fits (2.1) treating the imputed survival times as uncensored. These approaches, which we call VC : K and VC : M for “variance component,” use the same kernels as GPR : K and GPR : M, and are fit using the MM-algorithm from [Zhou and others \(2018\)](#).

Finally, we also consider two variations of the mixed-effects Cox proportional hazards model used in [Zhu and others \(2017\)](#). These versions are denoted ME-Cox : K and ME-Cox : M and use the same candidate kernels as GPR : K and GPR : M, respectively.

4.3. Performance metrics

As noted in [Van Wieringen and others \(2009\)](#), there is no consensus on which metric to use for evaluating the accuracy of prediction in survival analysis. For this reason, we use three different metrics. The first metric is based on the C-index measurement proposed by [Uno and others \(2011\)](#). Given $\hat{S}$, the n_v predicted survival times or risk scores for the testing set, we define the C-index as:

$$\frac{\sum_{i=1}^{n_v} \sum_{j=1}^{n_v} \delta_i \{\hat{H}(S_i)\}^{-2} 1(S_i < S_j, S_i < \tau) 1(\mathcal{C}(\hat{S}_i) > \mathcal{C}(\hat{S}_j))}{\sum_{i=1}^{n_v} \sum_{j=1}^{n_v} \delta_i \{\hat{H}(S_i)\}^{-2} 1(S_i < S_j, S_i < \tau)},$$

where $\mathcal{C} : \mathbb{R} \rightarrow \mathbb{R}$ maps a predicted survival time to the risk score scale, $1(\cdot)$ is the indicator function, τ is the study length, and $\hat{H}(\cdot)$ is the Kaplan–Meier estimator of the censoring distribution.

The second metric we use is the integrated Brier score. While C-index evaluates the prediction accuracy in terms of relative order of survival times, the Brier score quantifies the accuracy of survival time function estimate at time t :

$$B(t) = n_v^{-1} \sum_{i=1}^{n_v} \left\{ [\hat{S}(t|z_i, x_i)]^2 \min\{1(S_i \leq t), \delta_i\} \{\hat{H}(S_i)\}^{-1} + [1 - \hat{S}(t|z_i, x_i)]^2 1(S_i > t) \{\hat{H}(S_i)\}^{-1} \right\},$$

where $\hat{S}(t|z_i, x_i)$ is the estimated survival function for the i th subject in the testing data evaluated at time t . The integrated Brier score we use is

$$\text{BS} = \tau^{-1} \int_0^\tau B(t) dt.$$

The third metric we use is integrated AUC based on the sensitivity and specificity measures defined by [Uno and others \(2007\)](#). This AUC-based metric measures the accuracy of t -year survival prediction. We use the function `AUC.uno` from the R package `SurvAUC` in our implementation.

In the simulation study, we have access to the true survival times for the entire testing set, so $\hat{H}(Y_j) = 1$ and $\delta_j = 1$ for all j in the testing set. Thus, in the simulation study, we set τ equal to the largest survival time in the testing data. In the real data analyses, we estimate $\hat{H}$ using the Kaplan–Meier estimator and set τ equal to the second largest observed survival time in the testing data. For the methods of [Grimes and others \(2018\)](#), we compute survival probabilities using the inverse-probability-weighted approach from [Datta and others \(2007\)](#).

4.4. Results

Simulation results for 500 independent replications are displayed in Table 1. To compare methods, for each replication we record the ratio of each method's performance to the best performance amongst all methods. For C-index and integrated AUC, a ratio less than one indicates worse performance, whereas for integrated Brier score, a ratio greater than one indicates worse performance. For all metrics, a relative score closer to one indicates better relative performance.

Under Models 1–3, our proposed methods GPR:K and GPR:M were the best amongst the competing methods in terms of integrated Brier score. For integrated AUC and C-index, the version of our method with correctly specified covariance was best. Under Model 4, where the data were generated under the Cox proportional hazards model, our method still performs better than the methods developed under the Cox model in terms of AUC or C-index; but performs poorly in terms of integrated Brier score. This implies that the Cox model could provide more accurate estimates of the survival function, but less accurate prediction

Table 1. Average relative performance for 500 independent replications under the four models. Relative performance is defined as the ratio of each method's error to the best amongst all the competing methods, so that a relative performance of one indicates that the method performed best amongst all the methods. For integrated AUC and C-index, a relative error less than one indicates worse performance, whereas for integrated Brier score, a relative error of greater than one indicates worse performance. Cells highlighted in dark gray indicate a relative performance significantly better than all other methods. Cells highlighted in light gray indicate the best average relative performance, but one not significantly better than all other methods. The scale column displays the mean of the best method's metric across the 500 replications

Model	Covariance	Metric	Regularized AFT				Regularized Cox				ME-Cox		GPR		VC		Scale
			L_1	L_2	L_2^*	L_1^*	L_1	L_2	L_2^*	L_1^*	K	M	K	M	K	M	
1	Genome-wide	Brier	1.122	1.122	1.103	1.115	1.188	1.166	1.160	1.181	1.192	1.178	1.029	1.039	1.046	1.046	0.049
		AUC	0.969	0.975	0.958	0.952	0.965	0.975	0.977	0.966	0.946	0.944	0.989	0.984	0.967	0.966	0.795
		C-index	0.973	0.977	0.963	0.958	0.969	0.978	0.979	0.969	0.953	0.952	0.990	0.986	0.971	0.970	0.710
	Pathway	Brier	1.136	1.137	1.118	1.119	1.188	1.170	1.162	1.179	1.198	1.180	1.038	1.025	1.050	1.063	0.049
		AUC	0.961	0.965	0.973	0.969	0.955	0.965	0.968	0.958	0.933	0.938	0.979	0.992	0.969	0.957	0.802
		C-index	0.965	0.969	0.975	0.972	0.960	0.968	0.971	0.962	0.940	0.945	0.980	0.992	0.972	0.961	0.716
2	Genome-wide	Brier	1.109	1.111	1.093	1.105	1.181	1.158	1.151	1.170	1.182	1.168	1.026	1.035	1.048	1.046	0.047
		AUC	0.968	0.975	0.958	0.952	0.963	0.975	0.977	0.964	0.947	0.947	0.988	0.983	0.965	0.965	0.792
		C-index	0.972	0.978	0.963	0.958	0.967	0.978	0.979	0.968	0.953	0.953	0.989	0.984	0.969	0.969	0.709
	Pathway	Brier	1.134	1.138	1.117	1.116	1.190	1.172	1.166	1.182	1.195	1.180	1.039	1.026	1.054	1.065	0.048
		AUC	0.961	0.964	0.971	0.967	0.955	0.964	0.967	0.956	0.935	0.939	0.978	0.991	0.968	0.955	0.806
		C-index	0.964	0.968	0.974	0.971	0.959	0.968	0.970	0.960	0.942	0.946	0.979	0.992	0.970	0.960	0.720
3	Genome-wide	Brier	1.111	1.112	1.096	1.107	1.191	1.169	1.162	1.180	1.191	1.177	1.030	1.040	1.049	1.044	0.045
		AUC	0.969	0.974	0.960	0.955	0.964	0.975	0.977	0.965	0.944	0.943	0.989	0.983	0.965	0.967	0.793
		C-index	0.972	0.976	0.964	0.960	0.969	0.978	0.979	0.969	0.951	0.950	0.990	0.984	0.969	0.971	0.712
	Pathway	Brier	1.131	1.134	1.113	1.113	1.197	1.179	1.174	1.191	1.211	1.193	1.038	1.029	1.048	1.059	0.045
		AUC	0.962	0.966	0.973	0.969	0.958	0.967	0.969	0.958	0.936	0.939	0.981	0.992	0.970	0.960	0.804
		C-index	0.966	0.969	0.975	0.972	0.962	0.970	0.971	0.962	0.942	0.946	0.982	0.992	0.973	0.964	0.720
4	Genome-wide	Brier	1.082	1.089	1.070	1.074	1.068	1.055	1.059	1.075	1.080	1.072	1.160	1.167	1.085	1.079	0.092
		AUC	0.969	0.972	0.974	0.968	0.972	0.980	0.977	0.968	0.962	0.962	0.982	0.978	0.972	0.972	0.753
		C-index	0.974	0.976	0.978	0.973	0.977	0.983	0.981	0.973	0.968	0.969	0.984	0.982	0.976	0.976	0.680
	Pathway	Brier	1.093	1.095	1.078	1.079	1.072	1.058	1.062	1.081	1.086	1.078	1.152	1.149	1.078	1.077	0.092
		AUC	0.964	0.967	0.971	0.967	0.968	0.976	0.973	0.963	0.958	0.959	0.977	0.983	0.971	0.965	0.757
		C-index	0.970	0.972	0.976	0.973	0.973	0.980	0.977	0.969	0.965	0.966	0.981	0.986	0.976	0.971	0.682

of relative order. In all situations, our method performs better than VC, indicating that joint imputation and estimation generally yields better performance than doing imputation and estimation separately. Our method also performs better than the regularized AFT with L_1 and L_2 -penalties, which suggests that the kernel function captures the weak and abundant effects more effectively.

It is also important to analyze the performance of GPR : K when the true data generating covariance is based on pathway kernels. In general, GPR : K, which uses genome-wide gene expression, performs only slightly worse than GPR : M. Thus, even though only approximately 2.5% of genes actually contribute to the genomic effect under the pathway-based data generating model, using genome-wide gene expression does not seem to degrade prediction accuracy drastically.

In Figure 6(a) of the supplementary material available at *Biostatistics* online, we display boxplots showing the correlation between the true log-survival time and the imputed log-survival for the censored training data. The imputations compared are those obtained using the iterative method of Grimes and others (2018) based only on clinical/demographic variables, and those at convergence of our algorithm for both GPR : K and GPR : M. In summary, the correlations with the imputed log-survival times using our method were significantly higher than those obtained by the method of Grimes and others (2018), which would partly explain the relative advantage of GPR over VC in terms of prediction accuracy.

Because both our real data analysis and simulations were based on the KIRC data set with $n = 513$, we obtained a larger set of gene expression data from TCGA-BRCA ($n = 832$) and carried out an additional simulation study, detailed in Section 3.3 of the supplementary material available at *Biostatistics* online, to explore how the performance of our method varies as the training data sample size increases (from 432 to 732) with the testing set size fixed ($n = 100$). Briefly, the relative performances of the methods are consistent across training sample sizes. As one would expect, C-index and integrated AUC improve as the training sample size increases.

We also conducted additional simulation studies in Section 3.2 of the supplementary material available at *Biostatistics* online to examine the effects of a misspecified kernel. To summarize the results, we found using a misspecified kernel led to worse performance, though the difference across different kernel choices could be relatively small. While choosing the optimal kernel is a challenging problem, our method does provide a potential solution since we allow a practitioner to input M different types of kernels, and our method can estimate the relative contribution of each.

5. KIRC DATA ANALYSIS

5.1. Data preparation

We downloaded demographic/clinical data as well as RNA-seq data (from workflow HTSeq - Counts) of TCGA KIRC patients from NCI Genomic Data Commons (<https://portal.gdc.cancer.gov/>). We removed the patients with missing stage information or the patients for whom gene expression was measured on a plate with <10 patients, with 513 remaining for the following analysis. To filter out genes with low expression, we used genes whose 75 percentile read count was greater than 20. For the remaining 20 428 genes, the expression for the i th individual and k th gene is $x_{ik} = \log_{10}[(t_{ik} + 1)/q_{i,0.75}]$ where $q_{i,0.75}$ is the 75th percentile of read counts for the i th individual and t_{ik} is the read count for the i th individual's k th gene. We found that tumor stage, age, and sequencing plate had significant marginal associations with survival under the Cox proportional hazards model. Following Zhu and others (2017), we also include gender as a covariate.

5.2. Pathway-based candidate kernels

Following our simulation studies, we propose to use the normalized radial basis kernel $\underline{k}(\cdot, \cdot)$ to define the kernel function K . We consider multiple variations of our method: the genome-wide version, GPR : K,

which uses the normalized radial basis kernel based on genome-wide gene expression; and the pathway version, $\text{GPR} : \mathbf{M}$, which uses kernels computed using genes from individual pathways and a genome-wide gene expression kernel computed using all genes not included in any of the pathways.

We selected six pathways based on the existing knowledge of the molecular characteristics of KIRC. The PI3K/AKT/mTOR pathway (*PI3K*) was selected because it is recurrently mutated in KIRC patients (Cancer Genome Atlas Research Network, 2013). Since KIRC is characterized by a disordered metabolism, we also selected four pathways associated with metabolic function: glycolysis and gluconeogenesis (*Glyc/Gluc*), metabolism of fatty acids (*Fatty acid*), pentose phosphate pathway (*Pent/Phos*), and citrate cycle pathway (*Citrate cycle*). The genes belonging to three pathways *PI3K*, *Glyc/Gluc*, and *Fatty acid* were obtained from Molecular Signatures Database (<http://software.broadinstitute.org/gsea/msigdb>), and the genes of pathways *Pent/Phos* and *Citrate cycle* were obtained from Pathway Commons (<http://www.pathwaycommons.org/>). Finally, we included an immune-related gene-set, which includes genes that are expressed in different immune cell types and are used in the CIBERSORT software package (Newman and others, 2015) (*CSORT*). We included this gene set to capture potential immune-related information because previous studies have associated long survival time in KIRC patients with immune infiltration (Escudier, 2012).

We fit the multivariate regression of gene expression on sequencing plate, tumor stage, age, and gender, and compute the candidate kernels using the residuals. We found this adjustment had a minimal effect on survival time prediction compared with the analysis without adjustment.

5.3. Analyses and results

We randomly split the data into training and testing sets of size 413 and 100, respectively, for 500 independent replications. To compare the performance across methods, we use the same metrics defined in Section 4.3. Unlike our simulation studies, the test set contains censored survival times, so when computing the C-index and integrated Brier score, we estimate $\hat{H}$ using the Kaplan–Meier estimator on the testing data. The best C-index average was 0.746, the best integrated AUC average was 0.794, and the best integrated Brier score average was 0.155.

We display results for a subset of the competitors considered in our simulation studies in Figure 2. Both variations of our proposed method, $\text{GPR} : \mathbf{K}$ and $\text{GPR} : \mathbf{M}$, performed as well or better than the competitors across all three metrics. We noticed that the genome-wide version of our estimator tended to outperform the estimator which used pathway-based kernels. To further illustrate the contribution of each variance component in pathway-based analysis, we display boxplots of the relative magnitudes of the estimated variance components across the 500 replications (Figure 2(d)). We found that genome-wide effects accounted for approximately 60% of variability in log-survival time, whereas random noise accounted for approximately 20%. Of the considered pathways, the glycolysis and gluconeogenesis pathway accounted for 14% of variation, whereas the CIBERSORT genes account for approximately 6% of variability on average. Overall, these results suggest that most information are from genome-wide gene expression and the glycolysis and gluconeogenesis pathway also makes a significant contribution to survival time prediction.

6. DISCUSSION

In this article, we propose a new method to predict survival time using genome-wide gene expression data. Using the framework of Gaussian process regression, we develop a flexible and computationally efficient algorithm to perform two tasks: to impute censored survival time and to estimate the parameters of the model. We model the covariance structure of the log-survival time using one or more kernels defined using

Fig. 2. (a)–(c) Results from 500 independent training/testing splits. Each point represents the method’s performance relative to the best performing method within one replication. For (a) and (b), a relative C-index or relative integrated AUC less than one indicates worse performance. For (c), a relative Brier score of greater than one indicates a worse performance. (d) Estimated variance components for each of the candidate kernel in GPR : M.

gene expression data. In both simulations and real data analyses, we define multiple kernels using gene expression from multiple pathways, though in practice, these kernels can be defined using the same set of genes with different definitions of distance/similarities, or they can be defined based on multiple types of -omic data. Although we have developed our method for survival time prediction, it can be used or extended for other outcomes with certain patterns of missing data.

Our method was motivated by a data example where we found evidence of weak and abundant effects, and thus we do not advocate for using our method in settings with sparse and relatively strong effects, where variable selection is appropriate. However, additional simulations suggest that when the effects are sparse and relatively weak, our method still performs relatively well (see [Section 3.4 of the supplementary material](#) available at *Biostatistics* online).

Our model relies on a strong distributional assumption, which is needed partly for computational feasibility. In particular, even the truncated multivariate normal distribution requires intensive computation. We have conducted thorough simulations to evaluate three scenarios where our model assumptions are incorrect and conclude that our model can have better performance than the competing methods in all three scenarios. Therefore, we believe this strong assumption is a reasonable compromise for effective survival time prediction.

There are several directions to improve or extend our method. When the number of kernels is large, e.g., tens or hundreds of kernels, some regularization or penalty should be applied on the weight of each kernel. Although our simulations studies have demonstrated that our method is not sensitive to deviation from Gaussian distribution assumption, it is desirable to develop a more general parametric approach such as the Box–Cox transformation to allow for more flexibility than log-transformation, and to develop more robust non-parametric or semi-parametric approaches (Zeng and Lin, 2007).

SUPPLEMENTARY MATERIAL

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

ACKNOWLEDGMENTS

Conflict of Interest: None declared.

FUNDING

National Institutes of Health (NIH; R21CA224026, R01GM105785 R01CA189532).

REFERENCES

- BANERJEE, S., WALL, M. M., AND CARLIN, B. P. (2003). Frailty modeling for spatially correlated survival data, with application to infant mortality in Minnesota. *Biostatistics* **4**, 123–142.
- BENDER, R., AUGUSTIN, T. AND BLETTLNER, M. (2005). Generating survival times to simulate Cox proportional hazards models. *Statistics in Medicine* **24**, 1713–1723.
- CAFFO, B. S., JANK, W. AND JONES, G. L. (2005). Ascent-based Monte Carlo expectation–maximization. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67**, 235–251.
- CANCER GENOME ATLAS Research Network (2013). Comprehensive molecular characterization of clear cell renal cell carcinoma. *Nature* **499**, 43.
- DATTA, S. (2005). Estimating the mean life time using right censored data. *Statistical Methodology* **2**, 65–69.
- DATTA, S., LE-RADEMACHER, J. AND DATTA, S. (2007). Predicting patient survival from microarray data by accelerated failure time modeling using partial least squares and LASSO. *Biometrics* **63**, 259–271.
- DENG, D., DU, Y., JI, Z., RAO, K., WU, Z., ZHU, Y. AND COLEY, R. Y. (2016). Predicting survival time for metastatic castration resistant prostate cancer: an iterative imputation approach. *F1000Research* **5**, 2672–2672.
- ESCUDIER, B. (2012). Emerging immunotherapies for renal cell carcinoma. *Annals of Oncology* **23** (Suppl 8), viii35–viii40.
- FERNÁNDEZ, T., RIVERA, N. AND TEH, Y. W. (2016). Gaussian processes for survival analysis. In: *Proceedings of the 30th International Conference on Neural Information Processing Systems*. Barcelona, Spain: Curran Associates Inc., pp. 5021–5029.
- FLEGAL, J. M., HUGHES, J., VATS, D. AND DAI, N. (2017). *mcmcse: Monte Carlo Standard Errors for MCMC*. Riverside, CA, Denver, CO, Coventry, UK, and Minneapolis, MN. R package version 1.3-2. <https://cran.r-project.org/web/packages/mcmcse/index.html>.
- GÖNEN, M. AND ALPAYDIN, E. (2011). Multiple kernel learning algorithms. *Journal of Machine Learning Research* **12**, 2211–2268.
- GRIMES, T., WALKER, A. R., DATTA, S. AND DATTA, S. (2018). Predicting survival times for neuroblastoma patients using RNA-seq expression profiles. *Biology Direct* **13**, 11.

- HORRACE, W. C. (2005). Some results on the multivariate truncated normal distribution. *Journal of Multivariate Analysis* **94**, 209–221.
- KLEIN, J. P., PELZ, C. AND ZHANG, M.-J. (1999). Modeling random effects for censored data by a multivariate normal regression model. *Biometrics* **55**, 497–506.
- LIU, D., LIN, X. AND GHOSH, D. (2007). Semiparametric regression of multidimensional genetic pathway data: least-squares kernel machines and linear mixed models. *Biometrics* **63**, 1079–1088.
- NEWMAN, A. M., LIU, C. L., GREEN, M. R., GENTLES, A. J., FENG, W., XU, Y., HOANG, C. D., DIEHN, M. AND ALIZADEH, A. A. (2015). Robust enumeration of cell subsets from tissue expression profiles. *Nature Methods* **12**, 453.
- UNO, H., CAI, T., PENCINA, M. J., D’AGOSTINO, R. B. AND WEI, L. (2011). On the C-statistics for evaluating overall adequacy of risk prediction procedures with censored survival data. *Statistics in Medicine* **30**, 1105–1117.
- UNO, H., CAI, T., TIAN, L. AND WEI, L. (2007). Evaluating prediction rules for t-year survivors with censored regression models. *Journal of the American Statistical Association* **102**, 527–537.
- VAN WIERINGEN, W. N., KUN, D., HAMPEL, R. AND BOULESTEIX, A.-L. (2009). Survival prediction using gene expression data: a review and comparison. *Computational Statistics & Data Analysis* **53**, 1590–1603.
- WEI, G. C. AND TANNER, M. A. (1990). A Monte Carlo implementation of the EM algorithm and the poor man’s data augmentation algorithms. *Journal of the American Statistical Association* **85**, 699–704.
- WHEELER, H. E., AQUINO-MICHAELS, K., GAMAZON, E. R., TRUBETSKOY, V. V., DOLAN, M. E., HUANG, R. S., COX, N. J. AND IM, H. K. (2014). Poly-omic prediction of complex traits: OmicKriging. *Genetic Epidemiology* **38**, 402–415.
- WILHELM, S. AND MANJUNATH, B. (2015). *tmvtnorm: Truncated Multivariate Normal and Student t Distribution*. R Package Version 1.4-10. <https://cran.r-project.org/web/packages/tmvtnorm/index.html>.
- WU, C. F. J. (1983). On the convergence properties of the EM algorithm. *The Annals of Statistics* **11**, 95–103.
- WU, T., SUN, W., YUAN, S., CHEN, C.-H. AND LI, K.-C. (2008). A method for analyzing censored survival phenotype with gene expression data. *BMC Bioinformatics* **9**, 417.
- ZENG, D. AND LIN, D. (2007). Maximum likelihood estimation in semiparametric regression models with censored data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **69**, 507–564.
- ZHOU, H., HU, L., ZHOU, J. AND LANGE, K. (2018). MM algorithms for variance components models. *Journal of Computational and Graphical Statistics*, doi: 10.1080/10618600.2018.1529601.
- ZHU, B., SONG, N., SHEN, R., ARORA, A., MACHIELA, M. J., SONG, L., LANDI, M. T., GHOSH, D., CHATTERJEE, N., BALADANDAYUTHAPANI, V., AND ZHAO, H. (2017). Integrating clinical and multiple omics data for prognostic assessment across human cancers. *Scientific Reports* **7**, 16954.
- ZHU, R. AND KOSOROK, M. R. (2012). Recursively imputed survival trees. *Journal of the American Statistical Association* **107**, 331–340.

[Received August 24, 2018; revised April 7, 2019; accepted for publication May 13, 2019]