Published in final edited form as:

Stat Probab Lett. 2010 July 1; 80(13-14): 1056–1064. doi:10.1016/j.spl.2010.02.020.

Consistency of Random Survival Forests

Hemant Ishwaran and Udaya B. Kogalur

Cleveland Clinic

Abstract

We prove uniform consistency of Random Survival Forests (RSF), a newly introduced forest ensemble learner for analysis of right-censored survival data. Consistency is proven under general splitting rules, bootstrapping, and random selection of variables—that is, under true implementation of the methodology. Under this setting we show that the forest ensemble survival function converges uniformly to the true population survival function. To prove this result we make one key assumption regarding the feature space: we assume that all variables are factors. Doing so ensures that the feature space has finite cardinality and enables us to exploit counting process theory and the uniform consistency of the Kaplan-Meier survival function.

Keywords

Consistency; Ensemble; Factors; Kaplan-Meier; Random Forests; Survival Tree

1 Introduction

One of the most exciting machine learning algorithms to have been proposed is Random Forests (RF), an ensemble method introduced by Leo Breiman (Breiman, 2001). RF is an all-purpose algorithm that can be applied in a wide variety of data settings. In regression settings (i.e. where the response is continuous) the method is referred to as RF-R. In classification problems, or multiclass problems, where the response is a class label, the method is referred to as RF-C. Recently the methodology has also been extended to right-censoring survival settings, a method called random survival forests (RSF) (Ishwaran et al., 2008).

RF is considered an “ensemble learner”. Ensemble learners are predictors formed by aggregating many individual learners (base learners), each of which have been constructed from different realizations of the data. A widely used ensemble technique is bagging (Breiman, 1996). In bagging, the ensemble is formed by aggregating a base learner over independent bootstrap samples of the original data.

Although there are many variants of RF (Amit and Geman, 1997; Dietterich, 2000; Cutler and Zhao, 2001), the most popular, and the one we focus on here, is that described by Breiman in his software manual (Breiman, 2003) and discussed in Breiman (2001) under the name Forest-RI (short for RF random input selection). In this version, RF can be viewed as an extension of bagging. Using independent bootstrap samples, a random tree is grown by splitting each tree node using a randomly selecting subset of variables (features). The forest ensemble is

© 2009 Elsevier B.V. All rights reserved.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

constructed by aggregating over the random trees. The extra randomization introduced in the tree growing process is the crucial step distinguishing forests from bagging. Unlike bagging, each bootstrap tree is constructed using different variables, and not all variables are used. This is designed to encourage independence among trees, and unlike bagging, it not only reduces variance, but also bias. While this extra step might seem harmless, results using benchmark data have shown that prediction error for RF can be substantially better than bagging. In fact, performance of RF has been found comparable to other state-of-the-art methods such as boosting (Freund and Shapire, 1996), and support vector machines (Cortes and Vapnik, 1995).

Theoretical properties of RF have been investigated to some extent. In his seminal 2001 paper (Breiman, 2001), Breiman discussed bounds on the generalization error for a forest as a trade-off involving number of variables randomly selected as candidates for splitting, and the correlation between trees. He showed as number of variables increases, strength of a tree (accuracy) increases, but at the price of increasing correlation among trees; which degrades overall performance. In Lin and Jeon (2006), lower bounds for the mean-squared error for a regression forest were derived under random splitting by drawing analogies between forests and nearest neighbor classifiers. Recently, Meinshausen (2006) proved consistency of RF-R for quantile regression, and Biau, Devroye, and Lugosi (2008) proved consistency of RF-C under the assumption of random splitting.

In this paper, we prove consistency of RSF by showing that the forest ensemble survival function converges uniformly to the true population survival function. Because RSF is a new extension of RF to right-censored survival settings not much is known about its properties. Even consistency results for survival trees are sparse in the literature. For right-censored survival data, LeBlanc and Crowley (1993) showed survival tree cumulative hazard functions are consistent for smoothed cumulative hazard functions. The method of proof used convergence results for recursive-partitioned regression trees for uncensored data Breiman et.al (1984).

We take a different approach and establish consistency by drawing upon counting process theory. We first prove uniform consistency of survival trees, and from this, by making use of bootstrap theory, we prove consistency of RSF (Section 3). These results apply to general tree splitting rules (not just random ones) and to true implementations of RSF. We make only one important assumption: that the feature space is a finite (but very large) discrete space and that all variables are factors (Section 3). In this regard we deviate from other proofs of forest consistency which assume that the feature space is continuous. A continuous space is a more general assumption than ours and we readily acknowledge this limitation (see the Discussion). At that same time, discrete variables are very often encountered in medical settings involving survival data. Section 2.3 discusses an example related to esophageal cancer. Furthermore, in Section 4 we investigate the extent to which an assumption of a discrete feature space limits our results. We show by way of example that embedding forests in a discrete setting is realistic in that one can analyze problems with continuous variables by treating them as factors having a large number of factor labels. For the interested user, we note that all computations in the paper were implemented using the freely available R-software package, randomSurvivalForest (Ishwaran and Kogalur, 2007, 2008).

2 Random survival forests

Let $(\mathbf{X}, T, \delta)$, $(\mathbf{X}_1, T_1, \delta_1)$, $\dots$, $(\mathbf{X}_n, T_n, \delta_n)$ be i.i.d. random elements such that $\mathbf{X}$, the feature, is a d -dimensional vector taking values in $\mathcal{X}$, a discrete space (to be described in Section 3). Here $T = \min(T^0, C)$ is the observed survival time and $\delta = I(T^0 \leq C)$ is the binary $\{0, 1\}$ censoring value, where it is assumed that T^0 , the true event time, is independent of C , the censoring time.

An individual (case) i is said to be right-censored at time T_i if $\delta_i = 0$; otherwise, if $\delta_i = 1$, the individual is said to have experienced an event at T_i . It is assumed that $\mathbf{X}$ is independent of δ and that $(\mathbf{X}, T, \delta)$ has joint distribution $\mathbb{P}$. The marginal distribution for $\mathbf{X}$ is denoted by μ and defined via $\mu(A) = \mathbb{P}\{\mathbf{X} \in A\}$ for all subsets A of $\mathcal{X}$. It is assumed that $\mu(A) > 0$ for each $A \neq \emptyset$.

2.1 The RSF algorithm

The collection of values $\{(\mathbf{X}_i, T_i, \delta_i)\}_{1 \leq i \leq n}$ are referred to as the *learning data* and are used in the construction of the forest. We begin with a high-level description of the RSF algorithm. Specific details follow (Section 2.2).

1. Draw B independent Efron bootstrap samples (Efron, 1979) from the learning data and grow a binary recursive survival tree to each bootstrap sample.
2. When growing a survival tree, at each node of the tree randomly select p candidate variables to split on (use as many candidate variables as possible, up to $p - 1$, if there are less than p variables available within the node). Split the node by using the split that maximizes survival difference between daughter nodes (in the case of ties, a random tie breaking rule is used).
3. Grow a tree as near to saturation as possible (i.e. to full size) with the only constraint being that each terminal node should have no less than $d_0 > 0$ events.
4. Calculate the tree survival function. The forest ensemble is the averaged tree survival function.

2.2 Survival trees and forests: details of the algorithm

The tree survival function calculated in Step 4 of the algorithm is the Kaplan-Meier (KM) estimator for the tree's terminal nodes. To be more precise, let $\mathcal{N}(\mathcal{T})$ denote the terminal nodes of a survival tree, $\mathcal{T}$. These are the extreme nodes of $\mathcal{T}$ reached when the tree can no longer be split to form new nodes (daughters). Let $h \in \mathcal{N}(\mathcal{T})$ be a terminal node. Define $Y_h^{(n)}(t) = \sum_{i=1}^{m(h)} I(T_{i,h} \geq t)$ to be the number of individuals in h observed to be at risk just prior to time t (i.e., the number of individuals who have neither experienced an event nor been censored prior to t). Let $N_h^{(n)}(t)$ be the counting process defined as the number of events in $[0; t]$ for all cases in h . Define the indicator process $J_h^{(n)}(t) = I(Y_h^{(n)}(t) > 0)$. The Nelson-Aalen estimator for cases within the terminal node h is

$$\widehat{H}_h(t) = \int_0^t \frac{J_h^{(n)}(s)}{Y_h^{(n)}(s)} dN_h^{(n)}(s),$$

where we adopt the convention that $J_h^{(n)}(s)/Y_h^{(n)}(s) = 0$ whenever $Y_h^{(n)}(s) = 0$. The KM estimator for cases within h is

$$\widehat{S}_h(t) = \prod_{s \leq t} (1 - d\widehat{H}_h(s)) = \prod_{s \leq t} \left(1 - \frac{dN_h^{(n)}(s)}{Y_h^{(n)}(s)}\right). \quad (1)$$

Each case i has a d -dimensional feature $\mathbf{x}_i \in \mathcal{X}$. To determine the survival function for i , drop $\mathbf{x}_i$ down the tree. Because of the binary nature of a survival tree, $\mathbf{x}_i$ will be assigned a unique

terminal node $h' \in \mathcal{N}(\mathcal{T})$. The survival function for i is the KM estimator for $\mathbf{x}_i$'s terminal node:

$$\widehat{S}(t|\mathbf{x}_i) = \widehat{S}_{h'}(t), \text{ if } \mathbf{x}_i \in h'.$$

This defines the survival function for all cases and thus defines the survival function for the tree. To make this clear, we write the tree survival function as

$$\widehat{S}(t|\mathbf{x}_i) = \sum_{h \in \mathcal{N}(\mathcal{T})} I(\mathbf{x}_i \in h) \widehat{S}_h(t). \quad (2)$$

The forest ensemble survival function is the averaged tree survival function. If the forest is comprised of trees $\{\mathcal{T}_b\}_{1 \leq b \leq B}$, where $\widehat{S}_b(t|\mathbf{x})$ is the survival function for $\mathcal{T}_b$ given $\mathbf{x}$, the ensemble survival function is

$$\frac{1}{B} \sum_{b=1}^B \widehat{S}_b(t|\mathbf{x}) = \frac{1}{B} \sum_{b=1}^B \sum_{h \in \mathcal{N}(\mathcal{T}_b)} I(\mathbf{x}_i \in h) \widehat{S}_h(t).$$

Note that in practice each tree in the forest is constructed from an independent bootstrap sample of the data (Step 1 of the algorithm), but for now we ignore this complication. We revisit this issue later in Section 3.3 when we consider bootstrap resampling.

2.3 Esophageal cancer

Before delving into the theoretical results (Section 3) we present a previously published example illustrating RSF. In this example, RSF was applied to define a cancer stage grouping for esophageal cancer. Currently, staging of esophageal cancer is based solely on anatomic extent of disease using an orderly, progressive grouping of TNM cancer classifications. TNM stands for three anatomic features of esophageal cancer: depth of cancer invasion through the esophageal wall and adjacent tissues (T), presence of cancer-positive nodes along the esophagus (N), and presence of cancer metastases to distant sites (M). There are five subclassifications of T (Tis, T1, T2, T3, and T4), two of N (absence or presence of cancer-positive nodes), and two of M (absence or presence of distant metastases).

TNM classifications for esophageal cancer are known to be overly simplistic; in part because they reflect only anatomic extent of cancer. Other cancer characteristics known to affect prognosis include location of the cancer along the esophagus, cell type (squamous cell carcinoma vs. adenocarcinoma), and histologic grade; a crude reflector of biologic activity. It is also widely known that an increasing number of cancer-positive lymph nodes is associated with decreasing survival, and it is suspected this relationship is non-linear and may depend upon other factors, such as depth of cancer invasion, T.

In order to develop a more biologically plausible stage grouping, a RSF analysis was applied to a large group ($n = 4627$) of esophageal cancer patients (Ishwaran et al., 2009). Variables measured on each patient included TNM classifications, number of lymph nodes removed at surgery, number of cancer-positive lymph nodes, other non-TNM cancer characteristics, patient demographics, and additional variables to adjust for country and institution. All variables were discrete (this included patient age, which as customary was recorded in years). The primary outcome used in the analysis was time to death, measured from date of surgery.

Figure 1 displays some aspects of the forest constructed from this data. Plotted in the figure is five-year predicted survival for lymph node-positive patients who were free of distant metastases. The curves plotted are averaged ensemble survival functions evaluated at five years with each curve being constructed from a forest using a different number of trees (5, 10, 50, and 250 trees, respectively). Each curve is stratified by depth of cancer invasion (T1, T2, T3, and T4; note that Tis plays no role in node-positive cancers) and by number of cancer-positive nodes. Each point on a curve is the average of the ensemble survival function, averaged over all patients for a given T, for a given number of cancer-positive nodes. As the number of trees increases, the curves begin to stabilize, and after about 250 trees one can see there is not much to be gained by using additional trees. This shows that forest ensembles can stabilize fairly quickly. The other interesting aspect of Figure 1 is that RSF has clearly identified a non-linear relationship between survival and number of cancer-positive nodes, with survival decreasing rapidly as number of nodes increases. Furthermore, this relationship depends strongly upon T (the more deeply invasive the tumor, the less the effect). These results are consistent with the biology of esophageal cancer.

3 Properties of survival forests

3.1 Feature space

In establishing consistency of RSF we assume that each coordinate $1 \leq j \leq d$ of the d -dimensional feature $\mathbf{X}$ is a factor (discrete nominal variable) with $1 < L_j < \infty$ distinct labels. While this assumes that the feature space $\mathcal{X}$ has finite cardinality, the actual size of $\mathcal{X}$ can be quite large, $L_1 \times \cdots \times L_d$, and moreover, the number of splits that a tree might make from such data can be even larger, depending on d and L_j .

To see this, note that a split on a factor in a tree results in data points moving left and right of the parent node such that the complementary pairings define the new daughter nodes. For example, if a factor has three labels, $\{A, B, C\}$, then there are three complementary pairings (daughters) as follows: $\{A\}$ and $\{B, C\}$; $\{B\}$ and $\{C, A\}$; and $\{C\}$ and $\{A, B\}$. In general, for a factor with L_j distinct labels, there are $2^{L_j-1} - 1$ distinct complementary pairs. Thus, the total number of splits evaluated when splitting the root node for a survival tree when all variables are factors can be as much as

$$\text{maximum number root-node splits} = \sum_{j=1}^d 2^{L_j-1} - d.$$

Following the root-node split, are splits on the resulting daughter nodes, and their daughter nodes, recursively, with each subsequent generation requiring a large number of evaluations. Each evaluation can result in a new tree, thus showing that number of trees (space of trees) associated with $\mathcal{X}$ can be extremely large.

3.2 Uniform consistency of survival trees

To show that the ensemble survival function converges to the true population survival function we first prove consistency of a single tree. Consistency of forests will be readily deduced from this (Section 3.3).

In the following, and throughout the paper, the true survival function, or population parameter, is assumed to be of the form

$$S(t|\mathbf{X}) := \mathbb{P}\{T^o > t | \mathbf{X}\} = \sum_{\mathbf{x} \in \mathcal{X}} I(\mathbf{X} = \mathbf{x}) \exp\left(-\int_0^t \alpha(s|\mathbf{x}) ds\right), \quad (3)$$

where $\alpha(\cdot|\mathbf{x})$ is the non-negative hazard function for the subpopulation $\mathbf{X} = \mathbf{x}$.

The following result proves uniform consistency of a survival tree, and is a consequence of the uniform consistency of the KM estimator. The proof takes advantage of the finiteness of $\mathcal{X}$ which turns the problem of proving tree consistency into a more manageable problem of establishing consistency for a single terminal node.

Theorem 1. Let $t \in (0, \tau)$, where $\tau = \min\{\tau(\mathbf{x}) : \mathbf{x} \in \mathcal{X}\}$ and $\tau(\mathbf{x}) = \sup\{t : \int_0^t \alpha(s|\mathbf{x}) ds < \infty\}$. If $\mathbb{P}\{C > 0\} > 0$, and $\alpha(\cdot|\mathbf{x})$ is strictly positive over $[0, t]$ for at least one $\mathbf{x} \in \mathcal{X}$, then

$$\sup_{s \in [0, t]} \int_{\mathcal{X}} |\widehat{S}(s|\mathbf{x}) - S(s|\mathbf{x})| d\mu(\mathbf{x}) \xrightarrow{P} 0, \text{ as } n \rightarrow \infty,$$

where $\widehat{S}(\cdot|\mathbf{x})$ is a tree survival function defined as in (2).

3.3 Uniform consistency of survival forests

We have so far implicitly assumed that trees are grown from the learning data, $\mathcal{L} = \{(\mathbf{X}_i, T_i, \delta_i)\}_{1 \leq i \leq n}$. But in practice, trees are actually grown from independent bootstrap samples of the data. In order to prove consistency of a true implementation of RSF, we must extend our previous result to address bootstrap resampling.

Let $\mathcal{L}^* = \{(\mathbf{X}_i^*, T_i^*, \delta_i^*)\}_{1 \leq i \leq n}$ denote an Efron bootstrap sample (Efron, 1979) of $\mathcal{L}$. Let $\mathcal{T}^*$ be the survival tree grown from $\mathcal{L}^*$ and let $\widehat{S}^*(t|\mathbf{x})$ be the KM estimator for $\mathcal{T}^*$,

$$\widehat{S}^*(t|\mathbf{x}) = \sum_{h \in \mathcal{N}(\mathcal{T}^*)} I(\mathbf{x} \in h) \widehat{S}_h^*(t),$$

where $\widehat{S}_h^*(t)$ is defined similar to (1), and the above sum is over $\mathcal{N}(\mathcal{T}^*)$, the set of terminal nodes of $\mathcal{T}^*$.

The ensemble survival function for a survival forest comprised of B survival trees is

$$\widehat{S}_e(t|\mathbf{x}) := \frac{1}{B} \sum_{b=1}^B \widehat{S}_b^*(t|\mathbf{x}) = \frac{1}{B} \sum_{b=1}^B \sum_{h \in \mathcal{N}(\mathcal{T}_b^*)} I(\mathbf{x} \in h) \widehat{S}_h^*(t),$$

where $\widehat{S}_b^*(t|\mathbf{x})$ is the survival function for the survival tree $\mathcal{T}_b^*$ grown using the b th bootstrap sample. We prove consistency of RSF by establishing consistency of $\widehat{S}_b^*(t|\mathbf{x})$ for each b .

Theorem 2. Let $\tau^* = \min(\tau, \sup(F))$, where $\sup(F)$ is the upper limit of the support of $F(s) = 1 - \mathbb{P}\{T^o > s\} \mathbb{P}\{C > s\}$. Then under the same conditions as in Theorem 1, for each $t \in (0, \tau^*)$:

$$\sup_{s \in [0, t]} \int_{\mathcal{X}} |\widehat{S}^*(s|\mathbf{x}) - S(s|\mathbf{x})| d\mu(\mathbf{x}) = o_p^*(1) + o_p(1),$$

where o_p^* stands for o_p in bootstrap probability for almost all $\mathcal{L}$ -sample sequences; i.e. with probability one under $\mathbb{P}^\infty$.

Uniform consistency of the ensemble, $\hat{S}_e(t|\mathbf{x})$, follows automatically from Theorem 2, because

$$\begin{aligned} & \sup_{s \in [0, t]} \int_{\mathcal{X}} |\hat{S}_e(s|\mathbf{x}) - S(s|\mathbf{x})| d\mu(\mathbf{x}) \\ & \leq \frac{1}{B} \sum_{b=1}^B \sup_{s \in [0, t]} \int_{\mathcal{X}} |\hat{S}_b^*(s|\mathbf{x}) - S(s|\mathbf{x})| d\mu(\mathbf{x}). \end{aligned}$$

The right-hand side is $o_p^*(1) + o_p(1)$ by Theorem 2.

3.4 Uniform approximation by forests

Theorem 2 establishes consistency of a bootstrapped survival tree, and from this consistency of a survival forest follows. While this is a useful line of attack for establishing large sample properties of forests, it does not convey how in practice a forest might improve inference over a single tree. Indeed, in finite sample settings, a forest of trees can have a decided advantage when approximating the true survival function. Recall the esophageal cancer example of Section 2.3 where we saw first hand how averaging over trees improved prediction accuracy.

To provide theoretical support for superiority of forests, we make use of the following setting. Suppose that we are allowed to construct a binary survival tree $\mathcal{T}_b$ from a prechosen learning data set $\mathcal{L}_b = \{(\mathbf{X}_{b,i}, T_{b,i}, \delta_{b,i})\}_{1 \leq i \leq n}$ in any manner we choose. The only constraint being that each terminal node of $\mathcal{T}_b$ must contain at least $d_0 = 1$ events. We are allowed to construct $b = 1, \dots, B$ such trees for any $B < \infty$ of our choosing. Let $S_b(t|\mathbf{x})$ be the KM tree survival function for $\mathcal{T}_b$, and let

$$S_e(t|\mathbf{x}) = \sum_{b=1}^B W_b S_b(t|\mathbf{x}) \quad (4)$$

be the ensemble survival function constructed from $\{\mathcal{T}_b\}_{1 \leq b \leq B}$, where $\{W_b\}_{1 \leq b \leq B}$ are non-negative forest weights that we are free to specify. The next theorem shows that one can always find an ensemble that uniformly approximates the true survival function (3). Trees do not possess this property.

Theorem 3. *If $n > d$, and $s \in [0, \tau)$, then for each $\varepsilon > 0$ there exists an ensemble survival function (4) for a survival forest comprised of $B = B(\varepsilon)$ survival trees, with each tree consisting of $d + 1$ terminal nodes, such that*

$$\int_0^s \int_{\mathcal{X}} (S_e(t|\mathbf{x}) - S(t|\mathbf{x}))^2 \mu(d\mathbf{x}) dt \leq \varepsilon.$$

4 Treating a continuous problem as discrete

Our theory has been predicted on the assumption that all variables are factors, but in practice data with continuous variables are often encountered. Here we show that one can discretize continuous variables and treat them as factors without unduly affecting prediction error and inference: thus showing that our theory extrapolates reasonably to general data settings.

For illustration we consider the primary biliary cirrhosis (PBC) data of Fleming and Harrington (1991). The data is from a randomized clinical trial studying the effectiveness of the drug D-penicillamine on PBC. The dataset involves 312 individuals and contains 17 variables as well as censoring information and time until death for each individual. Of the 17 features, seven are discrete and 10 are continuous. Each of the 10 continuous variables were discretized and converted to a factor with L labels. We investigated different amounts of granularity: $L = 2, \dots, 30$.

For each level of granularity, L , we fit a survival forest of 1000 survival trees using log-rank splitting with node-adaptive random splits. Splits for nodes were implemented as follows. A maximum of “nsplit” complementary pairs were chosen randomly for each of the p randomly selected candidate variables within a node (if nsplit exceeded the number of cases in a node, then nsplit was set to the size of the node). Log-rank splitting was applied to the randomly selected complementary pairs and the node was split on that variable and complementary pair maximizing the log-rank test. Five different values for nsplit were tried: nsplit= 5, 10, 20, 50, 1024.

The top plot in Figure 2 shows out-of-bag prediction error as a function of granularity and nsplit value. As granularity rises, prediction error increases—but this increase is reasonably slow and well contained with larger values of nsplit. This is quite remarkable because total number of complementary pairs with a granularity level of $L = 30$ is on order 2^{30} (over 1 billion pairs) and yet our results show that using only a handful of randomly selected complementary pairs keeps prediction error in check.

Prediction error measures overall performance, but we should also consider how inference for a variable is affected by increasing granularity. To study this we looked at variable importance (VIMP). VIMP measures predictiveness of a variable, adjusting for all other variables (Ishwaran et al., 2008). Positive values of VIMP indicate predictiveness, and negative and zero values indicate noise. For each forest we dropped bootstrapped data down the forest and computed VIMP for each variable. This was repeated 1000 times independently. The bottom plot of Figure 2 displays the 68% bootstrap confidence region from this distribution. The analysis was restricted to only those forests grown under an nsplit value of 1024 but was carried out for each level of granularity (color coding scheme used to depict granularity is described in the caption of the figure). Overall, one can see that the bootstrap confidence regions are relatively robust to the level of granularity.

5 Discussion

We proved uniform consistency of RSF under settings that mirror those seen in actual data applications. Our consistency result followed by showing consistency of a single bootstrapped survival tree. As we remarked, while this is a useful line of attack for establishing large sample properties, it does not tell us why RF work. Theorem 3 presented one argument, but a general theory explaining why RF works still remains elusive. We hope that our work will motivate others to study this problem.

The major assumption in our approach was the assumption of a discrete feature space. Other proofs of consistency have used continuous feature spaces, and these results are more general. At the same time, because applications of RF to survival settings is much less studied than regression and classification settings, our results are useful because they add to the literature on forests. It is also interesting to note that a discrete space assumption implies, with probability one, that each node of the tree will contain one unique $\times$ value. This is in line with the original Forest-RI method suggested by Breiman (Breiman, 2001).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

- Amit Y, Geman D. Shape quantization and recognition with randomized trees. *Neural Computation* 1997;9:1545–1588.
- Andersen, PK.; Borgan, O.; Gill, RD.; Keiding, N. *Statistical Methods Based on Counting Processes*. New York: Springer; 1993.
- Biau G, Devroye L, Lugosi G. Consistency of random forests and other classifiers. *J. Machine Learning Research* 2008;9:2039–2057.
- Breiman, L.; Friedman, JH.; Olshen, RA.; Stone, CJ. *Classification and Regression Trees*. Belmont, California: 1984.
- Breiman L. Bagging predictors. *Machine Learning* 1996;26:123–140.
- Breiman L. Random forests. *Machine Learning* 2001;45:5–32.
- Breiman L. Manual – setting up, using, and understanding random forests V4.0. 2003 http://oz.berkeley.edu/users/breiman/Using_random_forests_v4.0.pdf.
- Cortes C, Vapnik VN. Support-vector networks. *Machine Learning* 1995;20:273–297.
- Cutler A, Zhao G. Pert - perfect random tree ensembles. *Computing Science and Statistics* 2001;33:490–497.
- Dietterich TG. An experimental comparison of three methods for constructing ensembles of decision trees: bagging, boosting, and randomization. *Machine Learning* 2000;40:139–157.
- Efron B. Bootstrap methods: another look at the jackknife. *Ann. Statist* 1979;7:1–26.
- Fleming, T.; Harrington, D. *Counting Processes and Survival Analysis*. New York: Wiley; 1991.
- Freund, Y.; Shapire, RE. Experiments with a new boosting algorithm; *Proc. of the 13th. Int. Conf. on Machine Learning*; 1996. p. 148-156.
- Ishwaran H, Kogalur UB. Random survival forests for R. *Rnews* 2007;7/2:25–31.
- Ishwaran H, Kogalur UB, Blackstone EH, Lauer MS. Random survival forests. *Ann. Appl. Stat* 2008;2:841–860.
- Ishwaran H, Blackstone EH, Hansen CA, Rice TW. A novel approach to cancer staging: application to esophageal cancer. *Biostatistics*. 2009 (in press).
- Ishwaran H, Kogalur UB. RandomSurvivalForest 3.5.0. R package. 2008 <http://cran.r-project.org>.
- LeBlanc M, Crowley J. Survival trees by goodness of split. *J. Amer. Stat. Assoc* 1993;88:457–467.
- Lin Y, Jeon Y. Random forests and adaptive nearest neighbors. *J. Amer. Statist. Assoc* 2006;101:578–590.
- Lo S-H, Singh K. The product-limit estimator and the bootstrap: some asymptotic representations. *Probab. Th. Rel. Fields* 1985;71:455–465.
- Meinshausen N. Quantile regression forests. *J. Machine Learning Research* 2006;7:983–999.

Figure 1.

RSF analysis of esophageal data. Five-year predicted survival for node-positive patients is plotted against number of cancer-positive nodes, stratified by depth of invasion (T1, T2, T3, and T4). Predicted survival is based on the forest comprising the first 5, 10, 50, and 250 trees, respectively.

Figure 2.

RSF analysis of PBC data using 1000 trees with random log-rank splitting where variables, both nominal and continuous, were discretized to have a maximum number of labels (factor granularity). Top figure is out-of-bag prediction error versus factor granularity, stratified by number of random splits used for a node, nsplit. Bottom figure shows 68% bootstrap confidence region for variable importance (VIMP) from 1000 bootstrap samples using an nsplit value of 1024 for each factor granularity value in the top figure. Color coding is such that the same color has been used for a variable over the different granularity values (factor granularity for a variable increases going from top to bottom).