Published in final edited form as:

Biometrics. 2020 December ; 76(4): 1177–1189. doi:10.1111/biom.13213.

ROC-guided survival trees and ensembles

Yifei Sun¹, Sy Han Chiou², Mei-Cheng Wang³

¹Department of Biostatistics, Columbia Mailman School of Public Health, New York, New York

²Department of Mathematical Sciences, University of Texas at Dallas, Richardson, Texas

³Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health, Baltimore, Maryland

Abstract

Tree-based methods are popular nonparametric tools in studying time-to-event outcomes. In this article, we introduce a novel framework for survival trees and ensembles, where the trees partition the dynamic survivor population and can handle time-dependent covariates. Using the idea of randomized tests, we develop generalized time-dependent receiver operating characteristic (ROC) curves for evaluating the performance of survival trees. The tree-building algorithm is guided by decision-theoretic criteria based on ROC, targeting specifically for prediction accuracy. To address the instability issue of a single tree, we propose a novel ensemble procedure based on averaging martingale estimating equations, which is different from existing methods that average the predicted survival or cumulative hazard functions from individual trees. Extensive simulation studies are conducted to examine the performance of the proposed methods. We apply the methods to a study on AIDS for illustration.

Keywords

concordance index; risk prediction; ROC curve; time-dependent covariate; tree-based method

1 | INTRODUCTION

Tree-based methods are popular alternatives to semiparametric and parametric methods. The basic idea of tree-based methods is to partition the covariate space into subsets (nodes), where individuals in the same node are alike regarding the outcome of interest. A single prediction is then assigned to individuals in the same node. In the setting of classification and regression trees (CART) (Breiman et al., 1984), this can be achieved by greedy splitting algorithms that minimize a measure of node impurity and the sum of squared deviations from the node mean, respectively. To improve predictive accuracy and reduce overfitting, one can apply a cost-complexity pruning algorithm to determine the size of the tree. In

Correspondence Yifei Sun, Department of Biostatistics, Columbia Mailman School of Public Health, New York, New York 10032. ys3072@cumc.columbia.edu.

SUPPORTING INFORMATION

The mathematical details and proofs referenced in Sections 2,3 and 7, the additional simulation results referenced in Section 4, and the additional data analysis referenced in Section 5 are available with this paper at the *Biometrics* website on Wiley Online Library.

practice, however, a small perturbation in the data may result in a substantial change in the structure of a fitted tree. Ensemble methods (Breiman, 1996; 2001) are ideal solutions to the instability problem and outperform a single tree in many applications.

With the increasing focus on personalized risk prediction, survival trees for time-to-event data have received much attention. A survival tree partitions the predictor space into a set of terminal nodes and reports the Kaplan–Meier estimate of the survival function in each node. In the literature, there has been a steady stream of works proposing new splitting rules to build survival trees. Gordon and Olshen (1985) first adopted the idea of CART and defined the node impurity to be the minimum Wasserstein distance between the Kaplan–Meier curves of the current node and a pure node. Alternatively, splitting rules that maximize the between node heterogeneity were commonly used in the literature: for example, Ciampi et al. (1986), Segal (1988), and LeBlanc and Crowley (1993) suggested selecting a split that yields the largest log-rank statistic; Moradian et al. (2017) used the integrated absolute difference between survival functions to measure the dissimilarity between child nodes. Moreover, researchers also considered likelihood-based splitting criteria, where the split is selected to maximize the sum of log-likelihoods from the two child nodes (LeBlanc and Crowley, 1992). Recently, the squared error loss in regression trees was extended to censored data (Molinario et al., 2004; Steingrimsen et al., 2016; 2019). Other splitting criteria include a weighted sum of impurity of the censoring indicator and the squared error loss of the observed event time (Zhang, 1995), and the Harrell’s C-statistic (Schmid et al., 2016). Readers are referred to Bou-Hamad et al. (2011b) for a comprehensive review of survival trees.

To address the instability issue of a single survival tree, researchers have developed various ensemble methods for time-to-event data. The basic algorithm operates by combining a large number of survival trees constructed from resampling the training data. For example, Hothorn et al. (2004) proposed a general method for bagging survival trees, and the final Kaplan–Meier estimate is computed using aggregated observations from all individual trees; Hothorn et al. (2006) proposed a random forest method to predict the log survival time; Ishwaran et al. (2008) proposed the random survival forest, where the estimates of the cumulative hazard function are averaged for the final prediction. Zhu and Kosorok (2012) used extremely randomized trees and proposed an imputation procedure that recursively updates the censored observations. Steingrimsen et al. (2019) considered more general weighted bootstrap procedures. Theoretical properties of random survival forests have been studied in Ishwaran and Kogalur (2010) and Cui et al. (2019).

The use of time-dependent covariates offers opportunities for exploring the association between failure events and risk factors that change over time. Applying standard methods such as the Cox proportional hazards model could be challenging in the choice of covariate form and can often yield biased estimation (Fisher and Lin, 1999). In contrast, tree-based methods allow the event risk to depend on the covariates flexibly. Bacchetti and Segal (1995) incorporated time-dependent covariates by using “pseudo-subject,” where the survival experience of one subject is viewed as experiences of multiple pseudo-subjects on nonoverlapping intervals, and survival probability using the truncation product-limit estimator is reported as the node summary. The idea of pseudo-subject was employed by

most of the existing works dealing with time-dependent covariates (Huang et al., 1998; Bou-Hamad et al., 2011a; Wallace, 2014; Fu and Simonoff, 2017). The pseudo-subject approach may have practical limitations, because one subject could be classified into multiple nodes in a tree, leading to a loss of simple interpretation and possible ambiguous prediction.

In this article, we propose a unified framework for survival trees and ensembles. To incorporate time-dependent covariates, we propose a time-invariant partition scheme on the survivor population. The partition-based risk prediction function is constructed using an algorithm guided by receiver operating characteristic (ROC) curves. Specifically, we define generalized time-dependent ROC curves and show that the target hazard function yields the largest area under the ROC curve. The optimality of the target hazard function motivates us to use a weighted average of the areas under the curve (AUC) on a time interval to guide splitting and pruning. Finally, we propose a novel ensemble method that averages unbiased martingale estimating equations instead of survival predictions from individual trees.

2 | SURVIVAL TREES WITH TIME-DEPENDENT COVARIATES

2.1 | A time-invariant partition on survivor population

Suppose T is a continuous survival time and $\mathbf{Z}(t)$ is a p -dimensional vector of possibly time-dependent covariates. Denote by $\lambda(t | \mathbf{z})$ the hazard function of T given $\mathbf{Z}(t) = \mathbf{z}$, that is,

$$\lambda(t | \mathbf{z}) dt = P\{T \in [t, t + dt) | \mathbf{Z}(t) = \mathbf{z}, T \geq t\}. \quad (1)$$

The function $\lambda(t | \mathbf{z})$ characterizes the instantaneous failure risk at t among survivors. At time t , let $\mathcal{X}_t$ denote the covariate space of $\mathbf{Z}(t)$ in the survivor population (i.e., the subpopulation satisfying $T \geq t$). To illustrate the idea, we first assume $\mathcal{X}_t = [0, 1]^p$ for $t \in (0, s]$ and s is a prespecified constant. We consider a partition on $\mathcal{X}_t$ that divides the survivor population into M groups, denoted by $\mathcal{T} = \{\tau_1, \dots, \tau_M\}$. The partition $\mathcal{T}$ is time-invariant and can be applied on $\mathcal{X}_t$ for all $t \in (0, s]$. The elements of the partition are disjoint subsets of $\mathcal{X}_t$ satisfying $\bigcup_{m=1}^M \tau_m = \mathcal{X}_t$ and are called terminal nodes. A subject enters a terminal node τ at t if $\mathbf{Z}(t) \in \tau$ and $T \geq t$. The partition $\mathcal{T}$ induces the following hazard model,

$$\lambda_{\mathcal{T}}(t | \mathbf{Z}(t)) = \sum_{\tau \in \mathcal{T}} I(\mathbf{Z}(t) \in \tau) \lambda(t | \tau), 0 < t \leq s, \quad (2)$$

where $\lambda(t | \tau) dt = P\{T \in [t, t + dt) | \mathbf{Z}(t) \in \tau, T \geq t\}$ is the node-specific hazard. Define a partition function $l_{\mathcal{T}}$ so that $l_{\mathcal{T}}\{\mathbf{z}\} = \tau$ if and only if $\mathbf{z} \in \tau, \tau \in \mathcal{T}$. The partition-based hazard function can then be written as $\lambda_{\mathcal{T}}(t | \mathbf{z}) = \lambda(t | l_{\mathcal{T}}\{\mathbf{z}\})$.

The time-invariant partition allows a sparse model and an easy interpretation of the decision rule. At each time t , the tree partitions the survivor population based on $\mathbf{Z}(t)$ and predicts the instantaneous failure risk. Thus, the interpretation at a fixed time point is along the same line as classification and regression trees. Since the risk within each terminal node changes with

time, it is essential to look at the hazard curves of each terminal node to determine the subgroups with high failure risks. Consider an example in Figure 1A where the partition $\mathcal{T} = \{\tau_1, \tau_2\}$ based on a single predictor (i.e., $p = 1$) divides the survivor population into two subgroups. Figures 1B,C are two possible scenarios the node-specific hazards can be defined. A larger value of the predictor is associated with a higher risk if the node-specific hazards are in Figure 1B, while a larger value of the predictor is associated with a lower risk in the early period and higher risk in the later period if the node-specific hazards are in Figure 1C. The partition-based hazard function $\lambda_{\mathcal{T}}(t | z)$ approximates the true hazard function $\lambda(t | z)$ as the partition becomes more refined.

Remark 1. The above discussion focuses on the case where $\mathcal{X}_t = [0, 1]^p$ for $t \in (0, s]$. In practice, if the domain of the time-dependent covariates among survivors changes over time, a time-invariant partition may not be appropriate. In this case, one can transform $\mathbf{Z}(t)$ onto $[0, 1]^p$ via a one-to-one function $G_t: \mathcal{X}_t \mapsto [0, 1]^p$. Let $\mathbf{X}(t) = G_t(\mathbf{Z}(t))$ be the transformed covariates and $h(t | \mathbf{x})$ be the hazard function of T given $\mathbf{X}(t) = \mathbf{x}$. A tree $\mathcal{T}'$ can be constructed using the transformed covariates $\mathbf{X}(t)$. Since $\lambda(t | z) = h(t | G_t(z))$, the tree-based hazard given $\mathbf{Z}(t) = z$ is $\lambda_{\mathcal{T}'}(t | G_t(z))$. One may use $G_t(z) = \hat{F}_t(z) \stackrel{\text{def}}{=} (\hat{F}_{1t}(z_1), \dots, \hat{F}_{pt}(z_p))$, where $z = (z_1, \dots, z_p)$ and for $q = 1, \dots, p$, $\hat{F}_{qt}$ is the empirical cumulative distribution function of the q th element of $\mathbf{Z}(t)$ among the at-risk subjects (i.e., $Y \geq t$).

2.2 | Estimation of tree-based risk function

Given a partition $\mathcal{T}$, we now consider the estimation of $\lambda_{\mathcal{T}}$ with right-censored survival data.

Let $Y = \min(T, C)$ be the observed survival time and $\delta = \mathbb{I}(T \leq C)$ be the failure event indicator. We use $\mathbf{Z}^H(t) = \{\mathbf{Z}(u), 0 \leq u \leq t\}$ to denote the covariate history up to t . The training data, denoted by $\mathcal{S}_n = \{Y_i, \Delta_i, \mathbf{Z}_i^H(Y_i); i = 1, \dots, n\}$, are independent identically distributed (i.i.d.) replicates of $\{Y, \mathbf{Z}^H(Y)\}$. To facilitate the discussion, we start from the case where there are a fixed number of terminal nodes in $\mathcal{T}$ and assume independent censoring within each terminal node, that is, $P\{T \in [t, t + dt) | \mathbf{Z}(t) \in \tau, T \leq t, C \leq t\} = \lambda(t | \tau)dt$ for $\tau \in \mathcal{T}$. As will be discussed later, a conditionally independent censoring assumption is imposed when the number of terminal nodes is allowed to increase with n . For $\tau \in \mathcal{T}$, define $F^*(t, \tau) = P\{Y \leq t, \delta = 1, \mathbf{Z}(Y) \in \tau\}$ and $S^*(t, \tau) = P\{\mathbf{Z}(t) \in \tau, Y \geq t\}$. Then we have $\lambda(t | \tau) = F^*(t, \tau)/S^*(t, \tau)$, where $F^*(t, \tau) = dF^*(t, \tau)/dt$. Define the observed counting process $N(t) = \mathbb{I}(Y \leq t)$. We estimate $F^*(t, \tau)$ by the following kernel type estimator,

$$\hat{f}^*(t, \tau) = \int_0^s K_h(t - u) d\hat{F}^*(u, \tau), t \in [h, s - h],$$

where $d\hat{F}^*(u, \tau) = \sum_{i=1}^n \mathbb{I}(\mathbf{Z}_i(Y_i) \in \tau) dN_i(u) / n$, $K_h(\cdot) = K(\cdot / h) / h$, $K(\cdot)$ is a second-order kernel function with a support on $[-1, 1]$, and h is the bandwidth parameter. To avoid biased estimation in the boundary region, one can either use the second-order boundary kernel (Müller, 1991), or set $\hat{f}^*(t, \tau) = \hat{f}^*(h, \tau)$ for $t \in [0, h)$ and $\hat{f}^*(t, \tau) = \hat{f}^*(s - h, \tau)$ for $t \in (s - h, s]$.

$s]$. Given a node τ , $\hat{f}^*(t, \tau)$ consistently estimates $f^*(t, \tau)$ as $n \rightarrow \infty$, $h \rightarrow 0$, and $nh \rightarrow \infty$. Note that $S^*(t, \tau)$ can be straightforwardly estimated by $\hat{S}^*(t, \tau) = \sum_{i=1}^n I(Z_i(t) \in \tau, Y_i \geq t) / n$. Thus the node-specific hazard $\lambda(t | \tau)$ can be estimated by

$$\hat{\lambda}(t | \tau) = \int_0^\infty K_h(t - u) d\hat{F}^*(u, \tau) / \hat{S}^*(t, \tau),$$

and $\lambda_{\mathcal{T}}(t | z)$ can be estimated by $\hat{\lambda}_{\mathcal{T}}(t | z) = \hat{\lambda}(t | l_{\mathcal{T}}\{z\})$.

In practice, the partition is usually constructed from data-dependent algorithms. To study the large sample property of the predicted hazard, we use $\mathcal{T}_n$ to denote a partition that can depend on the training data and define $\hat{\lambda}_{\mathcal{T}_n}(t | z) = \hat{\lambda}(t | l_{\mathcal{T}_n}\{z\})$. Given a new observation $Z_0(t)$ that is independent of the training data, we predict the hazard to be $\hat{\lambda}_{\mathcal{T}_n}(t | Z_0(t))$. We assume that the following conditions hold:

(A1) The censoring time C satisfies $P(t - T < t + dt | Z(t) = z, T \geq t, C \geq t) = \lambda(t | z)dt$.

(A2) The process $Z(t)$ is left-continuous and has right-hand limit. At time t , $Z(t)$ is distributed to a bounded density on $[0, 1]^p$. There exists a constant c_1 such that $f_{Z(t)|Y \geq t}(z)P(Y \geq t) > c_1$ for $z \in [0, 1]^p$, $t \in [0, s]$, where $f_{Z(t)|Y \geq t}(z)$ is the density of $Z(t)$ given $Y \geq t$.

(A3) The function $f^*(t | z) = \lim_{\delta_t \rightarrow 0^+} P\{T \in [t, t + \delta_t], \Delta = 1 | Z(Y) = z\} / \delta_t$ is

second-order differentiable with respect to t , and $\sup_{t \in [0, s], z \in [0, 1]^p} |f^*(t | z) / t^2| < c_2$ for some constant c_2 .

(A4) The number of terminal nodes in $\mathcal{T}_n$ grows at the rate of $\alpha(nh/\log n)$. For any $\gamma > 0$, the diameters of the nodes satisfies $\mu(z : \text{diam}(l_{\mathcal{T}_n}\{z\}) > \gamma) \rightarrow 0$ with probability 1, where μ is the Lebesgue measure.

(A5) The bandwidth h satisfies $h = n^{-a}$, $0 < a < 1$.

The conditional independent censoring assumption (A1) has been commonly adopted in survival analysis. Condition (A2) requires the support of $Z(t)$ in the at-risk population to be time-invariant on the time interval of interest. Condition (A3) requires the target function to be smooth so that kernel smoothing can be applied. Condition (A4) allows the size of the tree to become larger as the sample size increases. An analogous condition can be found in Nobel (1996), where the number of terminal nodes grows at the rate of $\alpha(n/\log n)$ for regression trees. The size of the tree is smaller in our case as our hazard estimator is based on kernel smoothing. The shrinking cell condition in (A4) is analogous to (12.9) in Breiman et al. (1984). Condition (A5) implies $h \rightarrow 0$ and $nh/\log n \rightarrow \infty$ as $n \rightarrow \infty$. We require the diameters of the terminal nodes and the bandwidth h to shrink toward zero so that the local estimate can approximate target function as $n \rightarrow \infty$. Similar to the existing works on tree-based estimators, the convergence result does not account for any specific splitting rule or

pruning procedure. Theorem 1 is developed to justify the hazard prediction under a set of regularity conditions. The proof is given in the Supporting Information.

Theorem 1. Under conditions (A1)–(A5), for $0 < t \leq s$ and any $\epsilon > 0$, as $n \rightarrow \infty$, we have $P\{|\hat{\lambda}_{\mathcal{T}_n}(t | \mathbf{Z}_0(t)) - \lambda(t | \mathbf{Z}_0(t))| > \epsilon | \mathcal{L}_n\} \rightarrow 0$ with probability 1.

Although the above discussion focuses on the hazard function, one can also predict the survival probability when $\mathbf{Z}(t)$ is a vector of *external* time-dependent covariates (Kalbfleisch and Prentice, 2011). Examples of external time-dependent covariates include fixed covariates, covariates determined in advance for each individual, and ancillary covariates whose probability laws do not involve parameters in the failure time model. Assume the hazard at t depends on $\mathbf{Z}^H(t)$ only through $\mathbf{Z}(t)$. The prediction of survival probability is based on the equation $P(T \geq t | \mathbf{Z}^H(t)) = \exp\{-\int_0^t \lambda(u | \mathbf{Z}(u)) du\}$. We predict the survival probability at t for a subject with covariate path $\mathbf{Z}_0^H(t) = \{\mathbf{Z}_0(u), 0 \leq u \leq t\}$ to be

$$\begin{aligned} \hat{P}(T \geq t | \mathbf{Z}_0^H(t)) \\ = \exp\left[-\int_0^t \frac{\sum_{i=1}^n I(\mathbf{Z}_i(u) \in l_{\mathcal{T}}\{\mathbf{Z}_0(u)\}) dN_i(u)}{\sum_{i=1}^n I(\mathbf{Z}_i(u) \in l_{\mathcal{T}}\{\mathbf{Z}_0(u)\}, Y_i \geq u)}\right]. \end{aligned} \quad (3)$$

3 | ROC-GUIDED SURVIVAL TREES

In this section, we propose an ROC-guided algorithm where the partition is constructed via a greedy approach that aims to maximize an ROC-related measure. The incident/dynamic time-dependent ROC curve (Heagerty and Zheng, 2005) is a popular tool for evaluating the prognostic accuracy of a continuous marker. Heuristically, at time t , the ROC curve is defined on the survivor population, where a subject is considered a case if $T = t$ and a control if $T > t$. Let $g(\cdot) : \mathcal{X}_t \mapsto \mathbb{R}$ be a scalar function that summarizes information from $\mathbf{Z}(t)$, and we predict $T = t$ or $T > t$ based on $g(\mathbf{Z}(t))$, with a larger value being more indicative of $T = t$. Following Heagerty and Zheng (2005), the false positive rate is $\text{FPR}_t(c) = P\{g(\mathbf{Z}(t)) > c | T > t\}$, the true positive rate is $\text{TPR}_t(c) = P\{g(\mathbf{Z}(t)) > c | T = t\}$, and the ROC function is $\text{ROC}_t(q) = \text{TPR}_t(\text{FPR}_t^{-1}(q))$. It is known that, when predicting a binary disease outcome with multiple disease markers, the risk score (i.e., the probability of disease given markers) yields the highest ROC curve (McIntosh and Pepe, 2002). For survival outcomes, the hazard $\lambda(t | \mathbf{Z}(t))$ can be viewed as an analog of the risk score. Following arguments of the Neyman–Pearson Lemma, setting $g(\cdot) = \lambda(t | \cdot)$ yields the highest ROC_t . Thus, $\hat{\lambda}_{\mathcal{T}_n}(t | \cdot)$ with a higher ROC_t curve is desired. However, when evaluating discrete-valued markers such as a tree-based risk score, the $\text{ROC}_t(\cdot)$ degenerates to a finite number of points. Hence, important summary measures such as AUC are not well defined. We fill in the gap by introducing a generalized time-dependent ROC curve.

3.1 | Generalized ROC curves for evaluating survival trees

With a finite number of terminal nodes, $\hat{\lambda}_{\mathcal{T}}(t | z)$ at a fixed time t is a discrete-valued scalar function of z . Thus, the ROC_t function for $\hat{\lambda}_{\mathcal{T}}(t | Z(t))$ becomes a finite set of points rather than a continuous curve. More generally, if $g(Z(t))$ has a point mass at c , the function $\text{ROC}_t(q)$ is undefined for $q \in (\text{FPR}_t(c), \text{FPR}_t(c-))$, where $\text{FPR}_t(c-) = \lim_{a \rightarrow c^-} \text{FPR}_t(a)$. In

this case, we construct a continuous curve, denoted by ROC_t^* , via linear interpolation.

Specifically, for $q \in (\text{FPR}_t(c), \text{FPR}_t(c-))$, the point $(q, \text{ROC}_t^*(q))$ on the ROC_t^* curve corresponds to the following prediction rule: if $g(Z(t)) > c$, predict $T = t$; if $g(Z(t)) = c$, predict $T = t$ with probability $\{q - \text{FPR}_t(c)\} / \{\text{FPR}_t(c-) - \text{FPR}_t(c)\}$; and if $g(Z(t)) < c$, predict $T > t$. In the special case where $g(Z(t))$ is a continuous variable, ROC_t^* reduces to ROC_t . The optimality of true hazard function $\lambda(t | \cdot)$ with respect to ROC_t^* is established in Proposition 1. Our result suggests that ROC_t^* can be used to evaluate the predictive ability of $\hat{\lambda}_{\mathcal{T}}(t | \cdot)$, and a higher ROC_t^* curve is favorable. The mathematical definition of ROC_t^* and the proof of Proposition 1 is given in the Supporting Information.

Proposition 1 (Optimality of the hazard function). *At time t , among all scalar functions $g: \mathcal{X}_t \mapsto \mathbb{R}$, the hazard $\lambda(t | \cdot)$ defined in (1) is optimal in the sense that $g(Z(t)) = \lambda(t | Z(t))$ yields the highest ROC_t^* curve.*

The area under the ROC_t^* curve, defined as $\text{AUC}_t^* = \int_0^1 \text{ROC}_t^*(q) dq$, has the interpretation of a concordance measure (Pepe, 2003). It can be shown that AUC_t^* is equivalent to

$$\begin{aligned} \text{CON}_t(g) &= P\{g(Z_1(t)) > g(Z_2(t)) \mid T_2 > T_1 = t\} \\ &\quad + \frac{1}{2}P\{g(Z_1(t)) = g(Z_2(t)) \mid T_2 > T_1 = t\}, \end{aligned}$$

where $\{Z_i(\cdot), T_i\}$, $i = 1, 2$, are i.i.d. replicates of $\{Z(\cdot), T\}$. Based on Proposition 1, CON_t can be used to evaluate the predictive ability of $g(Z(t))$. We then consider a global measure to evaluate $\hat{\lambda}_{\mathcal{T}}$ on $[0, s]$. Define a function $\tilde{g}: \mathbb{R}^+ \times \mathcal{X}_t \mapsto \mathbb{R}$ that combines $Z(t)$ in a time-dependent way and use $\tilde{g}(t, Z(t))$ to characterize the risk at t among survivors at t . We then integrate $\text{CON}_t(\tilde{g}(t, \cdot))$ over t with a weight function $\omega(t)$ and define,

$$\text{ICON}(\tilde{g}) = \int_0^s \omega(t) \text{CON}_t(\tilde{g}(t, \cdot)) dt.$$

In practice, the weight functions $\omega(t)$ can be specified by study investigators. A simple example is to set $\omega(\cdot) = 1$. Another possible choice is to set $\omega(t) = f(t)S(t)/P(T_2 > T_1, T_1 < s)$, where $f(t)$, $S(t)$ are the marginal density and survival functions of T , respectively. Then we have

$$\text{ICON}(\tilde{g}) = P\{\tilde{g}(T_1, \mathbf{Z}_1(T_1)) > \tilde{g}(T_1, \mathbf{Z}_2(T_1)) \mid T_2 > T_1, T_1 < s\} + 0.5P\{\tilde{g}(T_1, \mathbf{Z}_1(T_1)) = \tilde{g}(T_1, \mathbf{Z}_2(T_1)) \mid T_2 > T_1, T_1 < s\}$$

which is the probability that the subject failing earlier has a higher risk at the failure time. Following Proposition 1, the true hazard λ maximizes ICON. Motivated by this fact, we propose to use ICON as a guidance to build survival trees.

Remark 2. The Harrell's C-statistic (Harrell et al., 1982) has been commonly used to quantify the capacity of a risk score at baseline in discriminating among subjects with different event risks. When the event time is subject to censoring, the population parameter corresponding to the Harrell's C depends on the censoring distribution. Uno et al. (2011) studied a C-statistic that is consistent for a censoring free population concordance measure under a working Cox model $\lambda(t \mid \mathbf{X}) = \lambda_0(t) \exp(\mathbf{X}\boldsymbol{\beta})$, and $\mathbf{X}\boldsymbol{\beta}$ maximizes the limiting value of the Uno's C-statistic. However, without the Cox model assumption, it is not clear how to combine the predictors so that the limiting value is maximized. The proposed ICON is maximized when $\tilde{g}(t, \mathbf{z}) = \lambda(t \mid \mathbf{z})$ and is proper for guiding the tree building procedure.

3.2 | Splitting and pruning based on ICON

We next develop an algorithm to construct a time-invariant partition. Although the assumption (A1) is adopted for establishing the large-sample properties of the tree-based estimation, stronger assumptions are often generally to understand the splitting criteria, especially in the early steps of splitting. For example, when selecting the optimal split at the root node, the log-rank splitting rule implicitly assumes that C is independent of T within the child nodes, which is not guaranteed by assumption (A1). For ease of discussion, we assume C is independent of $\{T, \mathbf{Z}(\cdot)\}$ in Section 3.2. An extension to handle covariate-dependent censoring is given in Section 3 of the Supporting Information.

We first consider the estimation of $\text{ICON}(\lambda_{\mathcal{T}})$ using the training data. For a node τ , define $S(t, \tau) = P\{\mathbf{Z}(t) \in \tau, T > t\}$ and $f(t, \tau) = \lim_{\delta_t \rightarrow 0^+} P\{T \in [t, t + \delta_t), \mathbf{Z}(t) \in \tau\} / \delta_t$. Given a partition $\mathcal{T} = \{\tau_1, \tau_2, \dots, \tau_M\}$, the estimation of $\text{CON}_t(\lambda_{\mathcal{T}}(t \mid \cdot))$ is developed from

$$\text{CON}_t(\lambda_{\mathcal{T}}(t \mid \cdot)) = \frac{\sum_{j=1}^M \sum_{k=1}^M I\{\lambda(t \mid \tau_j) > \lambda(t \mid \tau_k)\} f(t, \tau_j) S(t, \tau_k) + 0.5 \sum_{j=1}^M f(t, \tau_j) S(t, \tau_j)}{\sum_{j=1}^M \sum_{k=1}^M f(t, \tau_j) S(t, \tau_k)}.$$

Under independent censoring, we have $f(t, \tau) = f^*(t, \tau)/P(C > t)$ and $S(t, \tau) = S^*(t, \tau)/P(C > t)$. Therefore, a consistent estimator for the concordance measure is given by

$$\begin{aligned} & \widehat{\text{CON}}_t(\hat{\lambda}_{\mathcal{T}}(t \mid \cdot)) \\ &= \frac{\sum_{j=1}^M \sum_{k=1}^M I\{\hat{\lambda}(t \mid \tau_j) > \hat{\lambda}(t \mid \tau_k)\} \hat{f}^*(t, \tau_j) \hat{S}^*(t, \tau_k) + 0.5 \sum_{j=1}^M \hat{f}^*(t, \tau_j) \hat{S}^*(t, \tau_j)}{\sum_{j=1}^M \sum_{k=1}^M \hat{f}^*(t, \tau_j) \hat{S}^*(t, \tau_k)} \quad (4) \end{aligned}$$

Note that the usual $O(n^2)$ computational costs for a concordance measure can be reduced by using (4), because the tree-based hazard $\lambda_{\mathcal{T}}(t \mid z)$ at time t takes M discrete values. To

estimate ICON, we use $\widehat{\text{ICON}}(\hat{\lambda}_{\mathcal{T}}) = \int_0^s \widehat{\text{CON}}_t(\hat{\lambda}_{\mathcal{T}}(t \mid \cdot)) \hat{w}(t) dt$, where $\hat{w}(t)$ is a weight function that possibly depends on the data. In practice, one can approximate the integral by the trapezoidal rule. Alternatively, one can also define ICON as a weighted average of CON_t on discrete time points. As demonstrated in our simulation studies, a moderate number of time points (e.g., 20) can yield reasonably good performances.

To build a partition, we begin at the top of the tree and then successively split the predictor space. At each splitting step, an optimal split can be chosen according to certain criteria. Consider the partition $\mathcal{T} = \{\tau_1, \dots, \tau_M\}$ and a split on any arbitrary node in $\mathcal{T}$. Without loss of generality, suppose τ_1 is split into τ_1^L and τ_1^R , and the partition after splitting is denoted by $\mathcal{T}' = \{\tau_1^L, \tau_1^R, \tau_m; m = 2, \dots, M\}$. Proposition 2 shows that splitting increases the concordance as long as the hazards of the two child nodes are different. The proof is given in the Supporting Information.

Proposition 2 (Splitting increases CON_t when two child nodes are heterogeneous). *Let $\lambda_{\mathcal{T}}$ and $\lambda_{\mathcal{T}'}$ be partition-based hazard functions before and after splitting, respectively. Then $\text{CON}_t(\lambda_{\mathcal{T}'}(t \mid \cdot)) \geq \text{CON}_t(\lambda_{\mathcal{T}}(t \mid \cdot))$, and the equality holds if and only if*

$$\begin{aligned} & \lambda(t \mid \tau_1) = \lambda(t \mid \tau_1^L) = \lambda(t \mid \tau_1^R). \text{ Moreover, we use “} \wedge \text{” to denote the estimated values using} \\ & (4), \text{ then } \widehat{\text{CON}}_t(\hat{\lambda}_{\mathcal{T}'}(t \mid \cdot)) \geq \widehat{\text{CON}}_t(\hat{\lambda}_{\mathcal{T}}(t \mid \cdot)), \text{ and the equality holds if and only if} \\ & \hat{\lambda}(t \mid \tau_1) = \hat{\lambda}(t \mid \tau_1^L) = \hat{\lambda}(t \mid \tau_1^R). \end{aligned}$$

Based on Proposition 2, we have $\widehat{\text{CON}}(\hat{\lambda}_{\mathcal{T}'}) \geq \widehat{\text{CON}}(\hat{\lambda}_{\mathcal{T}})$. When $\hat{w}(\cdot) > 0$, the equality holds if and only if $\hat{\lambda}(\cdot \mid \tau_1^L) = \hat{\lambda}(\cdot \mid \tau_1^R)$ almost everywhere on $(0, s]$. The validity of Proposition 2 does not depend on the censoring distribution. In practice, $\widehat{\text{CON}}(\hat{\lambda}_{\mathcal{T}'})$ may not correctly estimate $\text{ICON}(\lambda_{\mathcal{T}'})$ if the independent censoring assumption is violated, but the true concordance usually increases after splitting.

As ICON can effectively detect the difference in hazards of the two child nodes, it is natural to consider choosing the optimal split that maximizes ICON. However, an ICON-based

splitting rule is nonlocal because the split depends on the data in the parent node as well as other parts of the tree. As a result, additional computational burdens can arise when the number of terminal nodes in the current tree is large. In what follows, we introduce a local splitting rule, where the optimal split is chosen to maximize the increment of ICON within a node. For node τ 's child nodes τ^L and τ^R , it can be shown that CON_t within τ is 0.5 before splitting and is after splitting and is $0.5 + 0.5|f(t, \tau^L)S(t, \tau^R) - f(t, \tau^R)S(t, \tau^L)|f(t, \tau)S(t, \tau)$ after splitting. Hence, the increment in ICON within τ after splitting is

$$\Delta \text{ICON}_\tau = \int_0^s \frac{|f(t, \tau^L)S(t, \tau^R) - f(t, \tau^R)S(t, \tau^L)|}{f(t, \tau)S(t, \tau)} \omega(t) dt,$$

which can be estimated by

$$\begin{aligned} \widehat{\Delta \text{ICON}}_\tau &= \int_0^s \frac{|\hat{f}^*(t, \tau^L)\hat{S}^*(t, \tau^R) - \hat{f}^*(t, \tau^R)\hat{S}^*(t, \tau^L)|}{\hat{f}^*(t, \tau)\hat{S}^*(t, \tau)} \hat{\omega}(t) dt. \end{aligned}$$

Although splitting generally increases the concordance, a large tree can overfit the data, and the survival estimate within a small terminal node could be biased. Similar to the CART algorithm, we continue splitting until a predetermined stopping criterion on the minimal node size is met and then prune the fully grown tree. A node τ is considered to be

“splittable” only if $n(\tau) \stackrel{\text{def}}{=} \sum_{i=1}^n I(Z_i(0) \in \tau) \geq n_{\min}$, where $n_{\min}$ is a prespecified constant.

We also require all the nodes in the tree to satisfy that $n(\tau) \geq n_{\min}/2$. After the stopping criterion is met, we use the following concordance-complexity measure for pruning, where $\text{ICON}_\alpha(\mathcal{T}) = \widehat{\text{ICON}}(\hat{\lambda}_{\mathcal{T}}) - \alpha |\mathcal{T}|$, where $|\mathcal{T}|$ is the number of terminal nodes in $\mathcal{T}$ and α is a complexity parameter. For a fixed α , a tree with a large value of $\text{ICON}_\alpha(\mathcal{T})$ is generally favorable. Let K be the number of terminal nodes in the unpruned tree. For $k = 1, \dots, K$, let $\mathcal{T}_{(k)}$ be the size- k subtree that has the largest value of $\widehat{\text{ICON}}$. For each α , we define the optimal subtree $\mathcal{T}^\alpha$ as the subtree that has the largest ICON_α among $\mathcal{E} \stackrel{\text{def}}{=} \{\mathcal{T}_{(1)}, \dots, \mathcal{T}_{(K)}\}$.

For $\alpha_0 = 0$, the tree $\mathcal{T}^{\alpha_0} = \mathcal{T}_{(K)}$ is the optimal subtree. Define

$\alpha_{\mathcal{T}, \mathcal{T}'} = \{\widehat{\text{ICON}}(\hat{\lambda}_{\mathcal{T}}) - \widehat{\text{ICON}}(\hat{\lambda}_{\mathcal{T}'})\} / (|\mathcal{T}'| - |\mathcal{T}|)$. The q th ($q \geq 1$) threshold parameter

α_q is defined as $\alpha_q = \min\{\alpha_{\mathcal{T}, \mathcal{T}^{\alpha_{q-1}}}; |\mathcal{T}| \leq |\mathcal{T}^{\alpha_{q-1}}|, \mathcal{T} \in \mathcal{E}\}$ and $\mathcal{T}^{\alpha_q}$ is defined as the

smallest tree in $\{\mathcal{T} \mid \alpha_{\mathcal{T}, \mathcal{T}^{\alpha_{q-1}}} = \alpha_q, |\mathcal{T}| \leq |\mathcal{T}^{\alpha_{q-1}}|, \mathcal{T} \in \mathcal{E}\}$. Note that there exists an

integer Q such that $\mathcal{T}^{\alpha_Q} = \mathcal{T}_{(1)}$. For $\alpha \in [\alpha_q, \alpha_{q+1})$ and $q < Q$, $\mathcal{T}^{\alpha_q}$ is the optimal subtree;

and for $\alpha \in [\alpha_Q, \infty)$, $\mathcal{T}^{\alpha_Q}$ is the optimal subtree. In practice, α and $\mathcal{T}^{\alpha_Q}$ can be determined by cross-validation. The survival tree algorithm is given in Algorithm 1 outlined in Table 1.

Remark 3. In practice, one may consider node-specific bandwidths such that the bandwidth for node τ is $h_\tau = cn_\tau^{-1/5}$, where $n_\tau = \sum_{i=1}^n I(Z_i(Y_i) \in \tau)$. An order of $n_\tau^{-1/5}$ is chosen to achieve the lowest order of integrated mean square error within τ . Specifically, following the arguments of existing works on smoothing hazard, it can be shown that, $\int_0^s E\{\hat{\lambda}(t | \tau) - \lambda(t | \tau)\}^2 dt = O(h_\tau^4 + n_\tau^{-1} h_\tau^{-1}) = O(n_\tau^{-4/5})$. An ad hoc choice of c is $c_0 = s/8$ (Muller and Wang, 1994). In practice, one can also choose c via cross-validation.

4 | SURVIVAL ENSEMBLES BASED ON MARTINGALE ESTIMATING EQUATIONS

Survival trees can be transformed into powerful risk prediction tools by applying ensemble methods such as bagging (Breiman, 1996) and random forests (Breiman, 2001). The idea in bagging is to average many noisy but approximately unbiased tree models to reduce the variance; random forests further improve the variance by reducing the correlation between the trees via a random selection of predictors in the tree-growing process. In random forests for regression and classification (Breiman, 2001), the prediction for a new data point is the averaged prediction from individual trees that are often grown sufficiently deep to achieve low bias. For right-censored data, when the sizes of terminal nodes are very small, the within-node estimates of survival or cumulative hazard functions could be biased (Chen et al., 1982; Pena and Rohatgi, 1993). In what follows, we propose to average the unbiased martingale estimating equations rather than directly averaging node summaries. We treat forests as a type of adaptive nearest neighbor estimator (Lin and Jeon, 2006; Meinshausen, 2006; Athey et al., 2019) and propose local estimation for the survival or hazard functions.

Let $\mathbb{T} = \{\mathcal{T}_b\}_{b=1}^B$ be a collection of B partitions constructed using bootstrap samples. Each partition is constructed via a splitting procedure where at each split, m ($m < p$) predictors are randomly selected as candidates for splitting. The splitting criterion can either be the concordance measures in Section 3 or other appropriate criteria. Given $\mathcal{T}_b$ and a terminal node $\tau \in \mathcal{T}_b$, one can solve the following estimating equation for the node-specific hazard at t , $\sum_{i=1}^n w_{bi} I(Z_i(t) \in \tau) \{dN_i(t) - I(Y_i \geq t) \lambda(t | \tau) dt\} = 0$, where $w_{bi} \in \mathbb{N}_0$ is the frequency of the i th training observation in the b th bootstrap sample. Let $l_{\mathcal{T}_b}\{z\}$ be the partition function for $\mathcal{T}_b$ so that $l_{\mathcal{T}_b}\{z\} = \tau$ if and only if $z \in \tau$ and $\tau \in \mathcal{T}_b$, then the b th tree induces the following estimating equation for $\lambda(t | z)$,

$$\sum_{i=1}^n w_{bi} I(Z_i(t) \in l_{\mathcal{T}_b}\{z\}) \{dN_i(t) - I(Y_i \geq t) \lambda(t | z) dt\} = 0. \quad (5)$$

To get ensemble-based prediction, we take average of the estimating functions in (5) from all the B trees and obtain the following local martingale estimating equation for $\lambda(t | z)$,

$$\sum_{i=1}^n w_i(t, \mathbf{z}) \{dN_i(t) - I(Y_i \geq t) \lambda(t | \mathbf{z}) dt\} = 0, \quad (6)$$

where the weight function is $w_i(t, \mathbf{z}) = \sum_{b=1}^B w_{bi} I(\mathbf{Z}_i(t) \in l_{\mathcal{T}_b}(\mathbf{z})) / B$. The weight $w_i(t, \mathbf{z})$ captures the frequency with which the i th observation $\mathbf{Z}_i(t)$ falls into the same node as $\mathbf{z}$. Based on Equation (6), a kernel type estimator for the hazard function $\lambda(t | \mathbf{Z})$ is given by

$$\hat{\lambda}_{\mathbb{T}}(t | \mathbf{z}) = \int_0^\infty K_h(t - u) \frac{\sum_{i=1}^n w_i(u, \mathbf{z}) dN_i(u)}{\sum_{i=1}^n w_i(u, \mathbf{z}) I(Y_i \geq u)}. \quad (7)$$

Boundary correction is the same as that of survival trees. When $\mathbf{Z}(t)$ are external time-dependent covariates and the hazard at t depends on $\mathbf{Z}^H(t)$ only through $\mathbf{Z}(t)$, the survival probability at t given $\mathbf{Z}_0^H(t) = \{\mathbf{Z}_0(u), 0 \leq u \leq t\}$ is predicted to be

$$\exp \left\{ - \int_0^t \frac{\sum_{i=1}^n w_i(u, \mathbf{Z}_0(u)) dN_i(u)}{\sum_{i=1}^n w_i(u, \mathbf{Z}_0(u)) I(Y_i \geq u)} \right\}. \quad (8)$$

Details of the prediction procedure are given in Algorithm 2 outlined in Table 2.

The algorithm can also be extended to incorporate subsampling without replacement and sample-splitting (Athey et al., 2019). Specifically, one can divide the b th subsample from the original training data into two halves, whose indices are denoted by $\mathcal{T}_{1b}$ and $\mathcal{T}_{2b}$. Using the $\mathcal{T}_{1b}$ sample to place the splits and holding out the $\mathcal{T}_{2b}$ sample for within-leaf estimation yields honest trees. The honesty condition is proven to be successful in the literature on regression forests; readers are referred to Wager and Athey (2018) for an in-depth discussion. With subsampling and sample-splitting, the weight can be calculated as before but with w_{bi} being the frequency of the i th observation in $\mathcal{T}_{2b}$.

5 | SIMULATION STUDIES

We conducted extensive simulation studies to investigate the performance of the proposed methods. Our simulation settings are described below:

1. The survival times follow an exponential distribution with mean $\exp(0.1 \sum_{j=1}^{25} Z_j)$, where the covariate vector $(Z_1, \dots, Z_{25})$ follows a multivariate normal distribution with mean zero and a covariance matrix with the (i, j) th element equal to $0.9^{|i-j|}$. The censoring times follow an independent exponential distribution with mean η_1 .
2. The survival times follow an exponential distribution with mean $\sin(Z_1 \pi) + 2 | Z_2 - 0.5 | + Z_3^3$, where the elements in the covariate vector $(Z_1, \dots,$

Z_{25}) are i.i.d. uniform random variables on $[0, 1]$. The censoring times follow an independent uniform distribution on $[0, \eta_2]$.

3. The survival times followed a gamma distribution with shape parameter $0.5 + 0.3 \mid \sum_{j=1}^{15} Z_j \mid$ and scale parameter 2, where the covariate vector $(Z_1, \dots, Z_{25})$ follows a multivariate normal with mean zero and a covariance matrix with the (i, j) th element equal to $0.75^{|i-j|}$. The censoring times follow an independent uniform distribution on $[0, \eta_3]$.
4. The survival times follow a log-normal distribution with mean $\mu = 0.1 \mid \sum_{j=1}^5 Z_j \mid + 0.1 \mid \sum_{j=21}^{25} Z_j \mid$ and scale parameter one, where the covariate vector $(Z_1, \dots, Z_{25})$ is multivariate normal with mean zero and a covariance matrix with the (i, j) th element equal to $0.75^{|i-j|}$. The censoring times follow a log-normal distribution with mean $\mu + \eta_4$ and scale parameter one.
5. The hazard function is $\lambda(t \mid \mathbf{Z}(t)) = \exp\{0.5 \sum_{j=1}^{10} Z_j(t) + Z_{11}\} / 10$, where the time-dependent covariates $\{Z_1(t), \dots, Z_{10}(t)\}$ are multivariate normal with mean $kt + b$ and a covariance matrix with the (i, j) th element equal to $0.9^{|i-j|}$ at t , and the covariate $\{Z_{11}, \dots, Z_{20}\}$ is uniform on $[0, 1]$. The censoring times follow a uniform distribution on $[0, \eta_5]$.
6. The hazard function is $\lambda(t \mid \mathbf{Z}(t)) = \sum_{j=1}^{10} \{Z_j(t) - Z_{j+10}\}^2$. The time-dependent covariates are set as $Z_j(t) = tk_j/10$, with k_j drawn uniformly from $[0, 1]$. The covariates $(Z_{11}, \dots, Z_{20})$ are i.i.d. uniform random variables on $[0, 1]$. The censoring times follow a uniform distribution on $[0, \eta_6]$.
7. The hazard function is $\lambda(t \mid \mathbf{Z}(t)) = \sum_{j=1}^{10} \{Z_j(t) - Z_{j+10}\}^2$. The time-dependent covariates $Z_j(t) = tk_j/10$ and $(k_1, \dots, k_{10})$ are generated from a multivariate normal distribution with mean one and a covariance matrix with the (i, j) th element equal to $0.9^{|i-j|}$. The covariates $(Z_{11}, \dots, Z_{20})$ are i.i.d. standard normal random variables. The censoring times follow a uniform distribution on $[0, \eta_7]$.

Scenarios with only time-independent covariates were motivated from the settings in Zhu and Kosorok (2012) and Steingrimsen et al. (2019). The other three scenarios were included to examine the proposed methods in dealing with time-dependent covariates. For each setting, the censoring parameter η_j is tuned to yield censoring percentages of 0%, 25%, and 50%. The censoring distribution depends on covariates in Scenario 4. The proportional hazards assumption holds in Scenarios 1 and 5 but is violated in other scenarios.

We explored splitting based on the increment in within node ICON and overall ICON (i.e., $\Delta \widehat{\text{ICON}}_t$ in Algorithm 1 was replaced with the estimated ICON of the whole tree). All the covariates were transformed through $\hat{F}_t(\cdot)$ as in Remark 1 and treated as time-dependent covariates. For the ROC-guided survival trees, 10-fold cross-validation was used to choose the tuning parameter α in the concordance-complexity measure. For the ensemble method, 500 unpruned survival trees were constructed using bootstrap samples, and a set of randomly selected $\lceil \sqrt{p} \rceil$ features were considered at each split. For both methods, the minimum

number of baseline observations in each terminal node was set at 15. The Epanechnikov kernel function $K(x) = 0.75(1 - x^2)I(|x| \leq 1)$ was used, where the bandwidth $h = t_0/20$ for t_0 equals the 0.95 quantile of the uncensored survival times. The ICON measure was chosen as the average of CON_t on 20 equally spaced quantiles of the uncensored survival times.

The proposed methods were compared to several available implementations including the random survival forest (RSF) (Ishwaran and Kogalur, 2010), the relative risk tree (LeBlanc and Crowley, 1992), and the Cox proportional hazards model, implemented in the R packages randomForestSRC (Ishwaran and Kogalur, 2019), rpart (Therneau and Atkinson, 2018), and survival (Therneau, 2015), respectively. The RSFs were fitted using the default settings except that 500 trees were used in the ensembles. In addition to the default node size of 15, we also included an RSF algorithm using the node size that gives the smallest out-of-bag error. Tenfold cross-validation was used for pruning in relative risk trees. To the best of our knowledge, there are no existing software on survival trees and forests that can predict the survival probability based on covariate history, so both the RSF and the relative risk tree algorithms were fitted using the baseline values of all covariates.

The integrated absolute error (IAE) is used to evaluate the prediction accuracy of different methods. For each fitted model, we predicted the survival probabilities for $n_0 = 500$ new observations $\{Z_i^{new}(u), u \geq 0, i = 1, \dots, n_0\}$ generated from the distribution of the training data.

When there are only time-independent covariates, the IAE is defined as

$(n_0s)^{-1} \sum_{i=1}^{n_0} \int_0^s |\hat{P}\{T_i^{new} \geq t \mid Z_i^{new}(0)\} - P\{T_i^{new} \geq t \mid Z_i^{new}(0)\}| dt$, where s is the 0.95 quantile of the survival time. In the presence of time-dependent covariates, the IAE is defined as $(n_0s)^{-1} \sum_{i=1}^{n_0} \int_0^s |\hat{P}\{T_i^{new} \geq t \mid Z_i^{new,H}(t)\} - P\{T_i^{new} \geq t \mid Z_i^{new,H}(t)\}| dt$, where $Z_i^{new,H}(t) = \{Z_i^{new}(u), 0 \leq u \leq t\}$. For the proposed methods, the predicted survival probabilities were obtained from (3) and (8), with covariates transformed through $\hat{F}_t(\cdot)$.

Table 3 reports the summary statistics of the simulation results for sample sizes $n = 200$ and 1000. The results for $n = 100$ and 500, and the average terminal node sizes in ROC-guided trees after pruning are presented in the Supporting Information. The two splitting rules in the proposed methods yield similar IAEs for both trees and ensembles. Thus, we recommend the use of the $\widehat{\Delta CON}_t$ rule due to its lower computational cost. As expected, the proposed ensemble method has the lowest IAE in most scenarios, and the error decreases as n increases. Interestingly, with a small to moderate sample size (e.g., $n = 200$), our ensemble method outperforms the Cox model when simulating under the proportional hazard assumption of Scenarios 1 and 5. In the absence of time-dependent covariates, the proposed ensemble method and RSF generally perform better than single survival trees. On the other hand, the ROC-guided tree and the relative risk tree have similar performances. In the presence of time-dependent covariates, the proposed ensemble method continues to outperform all the other methods except for the Cox model when $n = 1000$ in Scenario 5. The ROC-guided tree also performs well but has larger IAEs than the Cox model in Scenario 5 when the proportional hazard assumption holds. The simulation results indicate that the

proposed methods perform well and incorporating time-dependent covariates improves the prediction accuracy substantially.

6 | APPLICATION

We illustrate the proposed methods through an application to a clinical trial conducted by Terry Bein Community Programs for Clinical Research on AIDS (Abrams et al., 1994; Goldman et al., 1996). The trial was conducted to compare didanosine (ddI) and zalcitabine (ddC) treatments for HIV-infected patients who were intolerant to or had failed zidovudine treatments. Of the 467 patients recruited for the study, 230 were randomized to receive the ddI treatment, and the other 237 received the ddC treatment. The average follow-up time was 15.6 months, and 188 patients died by the end of the study. Despite having longitudinal measurements that were measured at follow-up visits, Abrams et al. (1994) showed that the ddC treatment is more efficacious than the ddI treatment in prolonging survival time, based on a proportional hazards model with covariates measured at the baseline visit. We applied the proposed methods to investigate the time-dependent risk factors for overall survival. We included baseline covariates at randomization, such as gender, hemoglobin level, treatment received (ddI/ddC), and AIDS diagnosis (yes/no). We also included time-dependent covariates such as CD4 count, Karnofsky score, and cumulative recurrent opportunistic infections count. The CD4 count and Karnofsky score were measured at the baseline visit and bi-monthly follow-up visits. We adopted the last covariate carried forward approach between visit times when constructing these time-dependent covariates. For the opportunistic infection, we used the cumulative number of infections prior to t as the covariate value at t . As in Remark 1, the Karnofsky score and CD4 count were transformed into the range $[0,1]$ using the corresponding empirical cumulative distribution functions.

Figure 2A displays the proposed ROC-guided survival tree. After pruning, the terminal nodes are $\tau_1 = \{Z(t) \mid \text{KSC}(t) \leq 0.396\}$, $\tau_2 = \{Z(t) \mid \text{KSC}(t) > 0.396, \text{OP}(t) = 0\}$, and $\tau_3 = \{Z(t) \mid \text{KSC}(t) > 0.396, \text{OP}(t) > 0\}$, where $\text{KSC}(t)$ is the transformed Karnofsky score at t and $\text{OP}(t)$ is the cumulative number of opportunistic infection up to t . Here the transformed Karnofsky score of $\text{KSC}(t) = 0.396$ corresponds to Karnofsky scores between 65 and 75 depending on t . The partitions $\mathcal{T} = \{\tau_1, \tau_2, \tau_3\}$ corresponds to node 2, 6, and 7, whose estimated hazard rates are plotted in Figure 2B. Figure 2B clearly shows that a lower Karnofsky score is associated with higher mortality risk. For those with a high Karnofsky score, previous opportunistic infections are also associated with higher mortality risk.

Although treatment received is not selected as a splitting variable, the effect of the treatment on survival was predominantly mediated through the Karnofsky score and opportunistic infections. Thus, the survival tree provides insight on the mechanism by which the treatment operates. Moreover, the CD4 count is not used for splitting in the final tree, indicating that CD4 count may not be a surrogate endpoint in these patients; the result is consistent with the findings in Goldman et al. (1996). For comparison, the result from `rpart` that holds time-dependent covariates at baseline is given in the Supporting Information. `rpart` also makes a split using the Karnofsky score at the root node but yields a larger tree.

Finally, we applied the proposed ensemble algorithm. To visualize the hazard estimation from the survival ensemble, we plot the hazard functions over time for different Karnofsky scores and two cumulative opportunistic infection counts while holding all other covariates constant at the median (or the mode for binary covariates). The hazard curves in Figure 2C show that a low Karnofsky score (< 70 and dependent in normal activities) is associated with a higher risk of death, providing consistent results with our survival tree. On the other hand, for those with high Karnofsky scores (> 70 and independent in normal activities), the hazard estimates are fairly flat but can increase after the occurrence of opportunistic infection episodes.

7 | DISCUSSION

In this article, we propose a unified framework for survival trees and ensembles, where ROC_t^* and related summary measures guide the tree-growing algorithm. The proposed tree-based hazard estimators involve kernel smoothing and could result in biased estimation when the terminal nodes contain a small number of observations. However, with the proposed pruning procedure, very large trees with insufficient node sizes are not likely to be selected as the final model. In practice, if the true model is complex and a very large tree is needed to capture the underlying truth, we recommend the use of ensemble methods.

In the presence of time-dependent covariates, the node membership for the same subject could change over time. In practice, we usually evaluate partitions on a finite set of time points. The computational cost of our algorithm mainly depends on the sample size, the number of predictors, and the number of time points on which we evaluate the partitions. It has been shown that the computational cost of the CART algorithm is $O(pn \log(pn))$ (Sani et al., 2018). As is discussed in the Supporting Information, the computational cost for our algorithm that deals with time-dependent covariates is roughly $O(pqn \log(qn))$, where q is the number of time points on which CON_t is evaluated. Simulation studies showed $q = 20$ yields reasonably good performances when the sample size ranges from 100 to 1000.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

This work was supported by the National Institutes of Health grants U19 AG033655, R01 CA193888, R01 HL122212, and the Health Resources and Services Administration grant, T0BHP29302.

REFERENCES

- Abrams DI, Goldman AI, Launer C, Korvick JA, Neaton JD, Crane LR et al. (1994) A comparative trial of didanosine or zalcitabine after treatment with zidovudine in patients with human immunodeficiency virus infection. *New England Journal of Medicine*, 330, 657–662.
- Athey S, Tibshirani J and Wager S (2019) Generalized random forests. *The Annals of Statistics*, 47, 1148–1178.

- Bacchetti P and Segal MR (1995) Survival trees with time-dependent covariates: application to estimating changes in the incubation period of AIDS. *Lifetime Data Analysis*, 1, 35–47. [PubMed: 9385090]
- Bou-Hamad I, Larocque D and Ben-Ameur H (2011a) Discrete-time survival trees and forests with time-varying covariates: application to bankruptcy data. *Statistical Modelling*, 11, 429–446.
- Bou-Hamad I, Larocque D and Ben-Ameur H (2011b) A review of survival trees. *Statistics Surveys*, 5, 44–71.
- Breiman L (1996) Bagging predictors. *Machine learning*, 24, 123–140.
- Breiman L (2001) Random forests. *Machine Learning*, 45, 5–32.
- Breiman L, Friedman J, Stone CJ and Olshen RA (1984) *Classification and Regression Trees*. New York, NY: Chapman & Hall.
- Chen Y, Hollander M and Langberg N (1982) Small-sample results for the Kaplan-Meier estimator. *Journal of the American statistical Association*, 77, 141–144.
- Ciampi A, Thiffault J, Nakache J-P and Asselain B (1986) Stratification by stepwise regression, correspondence analysis and recursive partition: a comparison of three methods of analysis for survival data with covariates. *Computational Statistics & Data Analysis*, 4, 185–204.
- Cui Y, Zhu R, Zhou M and Kosorok M (2019) Consistency of survival tree and forest models: splitting bias and correction. [Preprint], Available at: arXiv:1707.09631.
- Fisher LD and Lin DY (1999) Time-dependent covariates in the Cox proportional-hazards regression model. *Annual Review of Public Health*, 20, 145–157.
- Fu W and Simonoff JS (2017) Survival trees for left-truncated and right-censored data, with application to time-varying covariate data. *Biostatistics*, 18, 352–369. [PubMed: 28025180]
- Goldman AI, Carlin BP, Crane LR, Launer C, Korvick JA, Deyton L and Abrams DI (1996) Response of CD4 lymphocytes and clinical consequences of treatment using ddI or ddC in patients with advanced HIV infection. *Journal of Acquired Immune Deficiency Syndromes*, 11, 161–169.
- Gordon L and Olshen RA (1985) Tree-structured survival analysis. *Cancer Treatment Reports*, 69, 1065–1069. [PubMed: 4042086]
- Harrell FE, Califf RM, Pryor DB, Lee KL and Rosati RA (1982) Evaluating the yield of medical tests. *Journal of the American Medical Association*, 247, 2543–2546. [PubMed: 7069920]
- Heagerty PJ and Zheng Y (2005) Survival model predictive accuracy and ROC curves. *Biometrics*, 61, 92–105. [PubMed: 15737082]
- Hothorn T, Bühlmann P, Dudoit S, Molinaro A and Van Der Laan MJ (2006) Survival ensembles. *Biostatistics*, 7, 355–373. [PubMed: 16344280]
- Hothorn T, Lausen B, Benner A and Radespiel-Tröger M (2004) Bagging survival trees. *Statistics in Medicine*, 23, 77–91. [PubMed: 14695641]
- Huang X, Chen S and Soong S-J (1998) Piecewise exponential survival trees with time-dependent covariates. *Biometrics*, 54, 1420–1433. [PubMed: 9883542]
- Ishwaran H and Kogalur U (2019) Fast unified random forests for survival, regression, and classification (RF-SRC). R package version 2.9.1. <https://cran.r-project.org/web/packages/randomForestSRC/index.html>.
- Ishwaran H and Kogalur UB (2010) Consistency of random survival forests. *Statistics & Probability Letters*, 80, 1056–1064. [PubMed: 20582150]
- Ishwaran H, Kogalur UB, Blackstone EH and Lauer MS (2008) Random survival forests. *The Annals of Applied Statistics*, 2, 841–860.
- Kalbfleisch JD and Prentice RL (2011) *The Statistical Analysis of Failure Time Data*. New York, NY: John Wiley & Sons.
- LeBlanc M and Crowley J (1992) Relative risk trees for censored survival data. *Biometrics*, 48, 411–425. [PubMed: 1637970]
- LeBlanc M and Crowley J (1993) Survival trees by goodness of split. *Journal of the American Statistical Association*, 88, 457–467.
- Lin Y and Jeon Y (2006) Random forests and adaptive nearest neighbors. *Journal of the American Statistical Association*, 101, 578–590.

- McIntosh MW and Pepe MS (2002) Combining several screening tests: optimality of the risk score. *Biometrics*, 58, 657–664. [PubMed: 12230001]
- Meinshausen N (2006) Quantile regression forests. *Journal of Machine Learning Research*, 7, 983–999.
- Molinaro AM, Dudoit S and Van der Laan MJ (2004) Tree-based multivariate regression and density estimation with right-censored data. *Journal of Multivariate Analysis*, 90, 154–177.
- Moradian H, Larocque D and Bellavance F (2017) L_1 splitting rules in survival forests. *Lifetime Data Analysis*, 23, 671–691. [PubMed: 27379423]
- Müller H-G (1991) Smooth optimum kernel estimators near endpoints. *Biometrika*, 78, 521–530.
- Muller H-G and Wang J-L (1994) Hazard rate estimation under random censoring with varying kernels and bandwidths. *Biometrics*, 50, 61–76. [PubMed: 8086616]
- Nobel A (1996) Histogram regression estimation using data-dependent partitions. *The Annals of Statistics*, 24, 1084–1105.
- Peña EA and Rohatgi VK (1993) Small sample and efficiency results for the Nelson-Aalen estimator. *Journal of Statistical Planning and Inference*, 37, 193–202.
- Pepe MS (2003) *The Statistical Evaluation of Medical Tests for Classification and Prediction*. Oxford: Oxford University Press.
- Sani HM, Lei C and Neagu D (2018) Computational complexity analysis of decision tree algorithms. In: Bramer M and Petridis M (Eds.) *International Conference on Innovative Techniques and Applications of Artificial Intelligence*. Berlin: Springer, pp. 191–197.
- Schmid M, Wright MN and Ziegler A (2016) On the use of Harrell's C for clinical risk prediction via random survival forests. *Expert Systems with Applications*, 63, 450–459.
- Segal MR (1988) Regression trees for censored data. *Biometrics*, 44, 35–47.
- Steingrimsson JA, Diao L, Molinaro AM and Strawderman RL (2016) Doubly robust survival trees. *Statistics in Medicine*, 35, 3595–3612. [PubMed: 27037609]
- Steingrimsson JA, Diao L and Strawderman RL (2019) Censoring unbiased regression trees and ensembles. *Journal of the American Statistical Association*, 114, 370–383. [PubMed: 31190691]
- Therneau T and Atkinson B (2018) rpart: recursive partitioning and regression trees. R package version 4.1-13. <https://cran.r-project.org/web/packages/rpart/index.html>.
- Therneau TM (2015) A package for survival analysis in S. version 2.38. <https://cran.r-project.org/web/packages/rpart/index.html>.
- Uno H, Cai T, Pencina MJ, D'Agostino RB and Wei LJ (2011) On the C-statistics for evaluating overall adequacy of risk prediction procedures with censored survival data. *Statistics in Medicine*, 30, 1105–1117. [PubMed: 21484848]
- Wager S and Athey S (2018) Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113, 1228–1242.
- Wallace M (2014) Time-dependent tree-structured survival analysis with unbiased variable selection through permutation tests. *Statistics in Medicine*, 33, 4790–4804. [PubMed: 25043382]
- Zhang H (1995) Splitting criteria in survival trees. *Statistical Modelling*, 104, 305–313.
- Zhu R and Kosorok MR (2012) Recursively imputed survival trees. *Journal of the American Statistical Association*, 107, 331–340. [PubMed: 23125470]

FIGURE 1.
Illustration of the survival tree and hazard prediction

(a) Partition on the survivors at time t . The nodes with * are terminal nodes.

(b) Hazard estimates for the groups represented by the terminal nodes 2, 6 and 7.

(c) Hazard estimates from the ensemble method given different Karnofsky scores, with $OP(t) = 0$ (top) and $OP(t) = 1$ (bottom).

FIGURE 2.

Mortality risks for the survivor population over time in the AIDS trial. This figure appears in color in the electronic version of this article, and any mention of color refers to that version

TABLE 1

The survival tree algorithm

Algorithm 1: The ROC-guided survival tree algorithm
Input : The training data are $\{Y_i, \Delta_i, Z_i^H(Y_i); i = 1, \dots, n\}$.
Output: A time-invariant partition that can be used for risk prediction.
Splitting Start from the root node, which is labeled as node 1. Let m denote the number of nodes (including both internal and terminal nodes) in the current tree and $m = 1$. Set $k = 1$.
while $k \leq m$ do
if node k is splittable then
Identify all possible splits on node k ;
Find the split that results in the largest $\widehat{\Delta\text{CON}}_\tau$;
Update the tree with the selected split and label the left and right children nodes to be node $m + 1$ and $m + 2$, respectively;
Set non-splittable children nodes as terminal nodes;
$m \leftarrow m + 2$;
end
$k \leftarrow k + 1$;
end
Pruning Calculate α_q and identify $\mathcal{T}^{\alpha_q}$ for $q = 1, \dots, Q$. For $q < Q$, set $\beta_q = \sqrt{\alpha_q \alpha_{q+1}}$ as the representative value of the interval $[\alpha_q, \alpha_{q+1})$ and $\beta_Q = \alpha_Q$;
Select the optimal β_q and the corresponding tree using cross-validation.

TABLE 2

The survival ensemble algorithm

Algorithm 2: The survival ensemble algorithm**Input :** The training data are $\{Y_i, \Delta_i, Z_i^H(Y_i); i = 1, \dots, n\}$.**Output:** The predicted hazard at t given z_t or the predicted survival probability at t given covariate history $\{z_u, 0 \leq u \leq t\}$.In what follows, define $|\mathcal{N}_b(u, z)| = \sum_{i=1}^n w_{bi} I(Y_i \geq u, Z_i(u) \in l_{\mathcal{T}_b}\{z\})$, where $w_{bi} \in \mathbb{N}_0$ is the frequency of the i th training observation in the b th bootstrap sample.**Hazard prediction**Initialize the weights: $(v_{01}, \dots, v_{0n}) = (v_{11}, \dots, v_{1n}) = (0, 0, \dots, 0)$.**for** $b = 1$ **to** B **do** Draw the b th bootstrap sample from the training data; Construct a partition $\mathcal{T}_b$ using the b th bootstrap sample with a random selection of m features at each split without pruning; **for** $i = 1$ **to** n **do** **if** $Z_i(Y_i) \in l_{\mathcal{T}_b}\{z_t\}$ **then** $v_{0i} \leftarrow v_{0i} + \Delta_i w_{bi}$; **end** $v_{1i} \leftarrow v_{1i} + |\mathcal{N}_b(Y_i, z_t)|$; **end****end**Predict the hazard at t given z_t using $\sum_{i=1}^n K_h(t - Y_i) v_{0i} / v_{1i}$.**Survival probability prediction**

Initialize the weights:

 $(v_{21}, v_{22}, \dots, v_{2n}) = (v_{31}, v_{32}, \dots, v_{3n}) = (0, 0, \dots, 0)$.**for** $b = 1$ **to** B **do** Draw the b th bootstrap sample from the training data; Construct a partition $\mathcal{T}_b$ using the b th bootstrap sample with a random selection of m features at each split without pruning; **for** $i = 1$ **to** n **do** **if** $Z_i(Y_i) \in l_{\mathcal{T}_b}\{z_{Y_i}\}$ **then** $v_{2i} \leftarrow v_{2i} + \Delta_i w_{bi}$; **end** $v_{3i} \leftarrow v_{3i} + |\mathcal{N}_b(Y_i, z_{Y_i})|$; **end****end**Predict the survival probability at t given $\{z_u, 0 \leq u \leq t\}$ using $\exp\{-\sum_{i=1}^n I(Y_i \leq t) v_{2i} / v_{3i}\}$.

Table 3:

Summaries of integrated absolute errors ($\times 1,000$). The numbers 0%, 25%, 50% correspond to the different censoring proportions. ROC-Tree and ROC-Ensemble are the proposed ROC-guided survival trees and ensembles with either the ICON or the ICON as the splitting criterion. RSF is the random survival forests methods implemented in R package randomForestSRC with either the default settings or with the nodesize parameter optimized based on out-of-bag error (using the function tune.nodesize). The Cox regression is implemented in R package survival. RR-Tree is the relative risk tree implemented in R package rpart.

See	ROC-Tree						ROC-Ensemble						RSF						RR-Tree					
	ICON			ICON			ICON			ICON			Default			optimal			Cox			0%		
	0%	25%	50%	0%	25%	50%	0%	25%	50%	0%	25%	50%	0%	25%	50%	0%	25%	50%	0%	25%	50%	0%	25%	50%
Scenarios with time-independent covariates																								
$n = 200$																								
I	82	95	152	81	95	151	49	54	89	49	54	89	59	70	126	74	87	134	48	57	99	87	101	146
II	103	107	166	104	108	165	77	79	124	77	79	124	75	79	130	90	101	161	110	118	175	82	84	131
III	121	127	175	120	127	176	101	103	135	100	103	134	99	102	136	119	131	175	152	158	196	122	125	160
IV	92	100	127	92	100	126	53	58	70	53	58	70	57	67	96	71	83	111	76	82	99	95	101	114
$n = 1000$																								
I	59	64	103	59	64	103	29	32	60	29	33	61	43	55	113	63	73	120	19	23	44	70	74	102
II	73	76	129	73	76	128	64	65	109	64	65	109	61	65	114	64	68	118	75	78	123	67	70	117
III	88	90	126	88	90	125	65	67	96	64	66	95	68	72	108	82	91	130	131	132	158	94	95	127
IV	68	71	83	67	71	83	35	38	46	35	38	46	41	52	81	53	59	79	34	38	46	74	78	89
Scenarios with time-dependent covariates																								
$n = 200$																								
V	77	84	104	77	84	105	67	71	86	67	71	86	151	140	138	152	139	135	71	75	91	167	149	126
VI	100	107	167	100	106	168	67	69	123	66	69	123	166	153	170	150	152	191	122	183	304	222	218	251
VII	116	119	179	116	119	179	91	94	146	91	94	146	150	144	176	160	167	215	131	165	242	190	188	221
$n = 1000$																								
V	64	65	74	64	65	75	57	57	69	56	57	69	152	137	133	155	138	134	59	59	64	158	139	122
VI	87	91	146	85	89	145	55	57	108	55	56	108	187	174	190	157	159	200	62	89	230	215	211	234
VII	108	112	165	108	115	170	84	85	132	83	84	132	162	153	179	162	174	227	92	103	162	182	174	202