



# An evaluation of machine learning techniques to predict the outcome of children treated for Hodgkin-Lymphoma on the AHOD0031 trial

Cédric Beaulac<sup>a</sup>, Jeffrey S. Rosenthal<sup>a</sup>, Qinglin Pei<sup>b</sup>, Debra Friedman<sup>c</sup>,  
Suzanne Wolden<sup>d</sup>, and David Hodgson<sup>e</sup>

<sup>a</sup>Department of Statistical Sciences, University of Toronto, Toronto, Canada; <sup>b</sup>Department of Biostatistics, University of Florida, Gainesville, Florida, USA; <sup>c</sup>Department of Pediatrics, Vanderbilt University, Nashville, Tennessee, USA; <sup>d</sup>Department of Radiation Oncology, Memorial Sloan Kettering Cancer Center, New York, New York, USA; <sup>e</sup>Department of Radiation Oncology, University of Toronto, Toronto, Canada

## ABSTRACT

In this manuscript, we analyze a data set containing information on children with Hodgkin Lymphoma (HL) enrolled on a clinical trial. Treatments received and survival status were collected together with other covariates such as demographics and clinical measurements. Our main task is to explore the potential of machine learning (ML) algorithms in a survival analysis context in order to improve over the Cox Proportional Hazard (CoxPH) model. We discuss the weaknesses of the CoxPH model we would like to improve upon and then we introduce multiple algorithms, from well-established ones to state-of-the-art models, that solve these issues. We then compare every model according to the concordance index and the Brier score. Finally, we produce a series of recommendations, based on our experience, for practitioners that would like to benefit from the recent advances in artificial intelligence.

## Introduction

There is increasing effort in medical research to applying ML algorithms to improve treatment decisions and predict patient outcomes. In this article, we want to explore the potential of ML algorithms to predict the outcome of children treated for Hodgkin Lymphoma. As we want to minimize the side effects of intensive chemotherapy or radiation therapy, a major clinical concern is how, for a given patient, we can select a treatment that eradicates the disease while keeping the intensity of the treatment, and the associated side effects, to a minimum.

In this article, we will introduce multiple ML algorithms adapted to our needs and compare them with the Cox proportional hazard model. As it is the case with many data set within this field, the response variable, time until death

or relapse, was right-censored for patients without events and the data set is of relatively small size ( $n = 1712$ ). From a ML perspective, this can be challenging. The response variable is right-censored for multiple observations but many ML techniques are not designed to deal with censored observations and thus it restricts the techniques we can include in our case study. Another challenge previously mentioned is that medical data sets are usually smaller than those used in ML applications and thus we will have to carefully select algorithms that could perform well in this context.

We will introduce the data set in Section 2. In Section 3, we will introduce the algorithms tested. Then, in Section 4, we will present our experimental setup and our results. Finally, in Section 5, we will discuss thoroughly the results, recommend further improvements and introduce open questions.

## Data Set

We have a data set of 1,712 patients, treated on the Children's Oncology Group trial AHOD0031, the largest randomized trial of pediatric HL ever conducted. Each observation represents a patient suffering from Hodgkin Lymphoma. For every patient, characteristics and symptoms have been collected as well as the treatment, for a total of 21 predictors. A table containing information on the predictors is in the Appendix. The response is a time-to-event variable registered in the number of days. We consider events to be either death or relapse. For patients without events, the response variable was right-censored at the time of last seen, which is a well-known data structure in survival analysis. This data set and the data collecting technique are presented in detail by Friedman et al. (2014) who previously analyzed the same data set for other purposes.

## Survival Analysis Models

### ***Benchmark: Cox Proportional Hazard Model***

The Cox Proportional Hazard (CoxPH) model (Cox 1972) serves as our benchmark model. It is widely used in medical sciences since it is robust, easy to use, and produces highly interpretable results. It is a semi-parametric model that fits the hazard function, which represents the instantaneous rate of occurrence for the event of interest, using a partial likelihood function (Cox 1975).

The CoxPH model fits the hazard function which contains two parts, a baseline hazard function of the time and a feature component which is a linear function of the predictors. The proportional hazard assumption assumes the time component and the feature component of the hazard function are proportional. In other words, the effect of the features is fixed through time.

In the CoxPH model, the baseline hazard, which contains the time component, is usually unspecified so we cannot use the model directly to compute the hazard or to predict the survival function for a given set of covariates.

The main goal of this analysis is to test whether or not new ML models can outperform the CoxPH model. As ML models have shown great potential in many data analysis applications, it is important to test their potential to improve outcome prediction for cancer patients. We would like our selected models to improve upon at least one of the three following problems that are intrinsic to the CoxPH model. Problem (1): the proportional hazard assumption; we would like models that allow for feature effects to vary through time. Problem (2): the unspecified baseline hazard function; we would like models able to predict the survival function itself. Problem (3): the linear combination of features; we would like to use models that are able to grasp high order of interaction between the variable or non-linear combinations of the features.

## ***Conventional Statistical Learning Models***

### ***Regression Models***

The first model to be tested is a member of the CoxPH family. One way to capture interactions between predictors in linear models, and thus improve toward problem (3), is to include interaction terms. Since typical medical data sets contain few observations and many predictors, including all interactions usually lead to model saturation.

To deal with this issue we will use a variable selection model. Cox-Net (Simon et al. 2011) is an extension of the now well-know lasso regression (Hastie, Tibshirani, and Friedman 2009) implemented in the glmnet package (Friedman, Hastie, and Tibshirani 2010) and is the first model we will experiment with. The Cox-Net is a lasso regression-style model that shrinks some model coefficients to zero and thus insures the model is not saturated. The resulting model is as interpretable as the benchmark CoxPH model, but Cox-Net allows us to include all interactions in the base model without losing too many degrees of freedom.

Another approach based on regression models is the Multi-Task Logistic Regression (MTLR). Yu et al. (2011) proposed the MTLR model which quickly became a benchmark in the ML community for survival analysis and was cited by many authors (Fotso 2018; Jinga et al. 2019; Luck et al. 2017; Zhao and Feng 2019). The proposed technique directly models the survival function by combining multiple local logistic regression models and considers the dependency of these models. By modeling the survival distribution with a sequence of dependent logistic regression, this model captures the time-varying effects of features and thus the proportional hazard assumption is not needed. The model also grants the ability to predict survival time for individual patients. This model solves both problems (1) and (2). For our case study, we used the MTLR R-package (Haider 2019) recently implemented by Haider.

### ***Survival Tree Models***

Decision trees (Breiman et al. 1984) and random forests (Breiman 1996, 2001) are known for their ability to detect and naturally incorporate high degrees of interactions among the predictors which is helpful toward problem (3). This family of models is well established and makes very few assumptions about the data set, making it a natural choice for our case study.

Multiple adaptations of decision trees were suggested for survival analysis and are commonly referred as survival trees. The idea suggested by many authors is to modify the splitting criteria of decision trees to accommodate for right-censored data. Based on previously published reviews of survival trees (Bou-Hamad, Larocque, and Ben-Ameur 2011; LeBlanc and Crowley 1995), we have selected four techniques for the case study.

One of the oldest survival tree models that was implemented in R (R Core Team 2013) is the Relative Risk Survival Tree (Leblanc and Crowley 1992). This survival tree algorithm uses most of the architecture established by CART (Breiman et al. 1984) but also borrows ideas from the CoxPH model. The model suggested by LeBlanc et al. assumes proportional hazards and partitions the data to maximize the difference in relative risk between regions. This technique was implemented in the `rpart` R-package (Therneau, Atkinson, and Ripley 2017).

We also selected a few ensemble methods. To begin, Hothorn et al. (2004) proposed a new technique to aggregate survival decision trees that can produce conditional survival function, which solves problem (2). To predict the survival probabilities of a new observation, they use an ensemble of survival trees (Leblanc and Crowley 1992) to determine a set of observations similar to the one in need of a prediction. They then use this set of observations to generate the Kaplan–Meier estimates for the new one. Their proposed technique is available in the `ipred` R-package (Peters and Hothorn 2019). A year later, Hothorn et al. (2005,) jointly with Strobl et al. (2007) proposed a new ensemble technique able to produce log-survival time estimates instead. We will test this technique that is implemented in `party` R-package (Hothorn et al. 2019; Hothorn, Hornik, and Zeileis 2006).

Finally, the latest development in random forests for survival analysis is Random Survival Forests (Ishwaran et al. 2008). This implementation of a random survival forest was shown to be consistent (Ishwaran and Kogalur 2010) and it comes with high-dimensional variable selection tools (Ishwaran et al. 2010). This model was implemented in the `randomForestSRC` R-package (Ishwaran and Kogalur 2019).

### ***State-of-the-art Models***

#### ***Deep-learning Models***

The first state-of-the-art model we will experiment with is built upon the most popular architecture of models in recent years: deep neural networks. Yu et al.

(2011) MTLR model inspired many modifications (Fotso 2018; Jinga et al. 2019; Luck et al. 2017; Zhao and Feng 2019) in order to include a deep-learning component to the model. The main purpose is to allow for interactions and non-linear effects of the predictors. Fotso (2018, 2019) suggested an extension of the MTLR where a deep neural network parameterization replaces the linear parameterization, and Luck et al. (2017) proposed a neural network model that produces two outputs: one is the risk and one is the probability of observing an event in a given time bin. Unfortunately, the authors for most of these techniques (Jinga et al. 2019; Luck et al. 2017; Zhao and Feng 2019) did not provide either their code or a package which causes great reproducibility problems and leads to a serious accessibility issue for practitioners. The DeepSurv architecture (Katzman et al. 2018) proposed by Katzman et al. is a direct extension to the CoxPH model where the linear function of the covariance is replaced by a deep neural network. This allows the model to grasp high order of interactions between predictors therefore solving problem (3). By allowing for interaction between covariates and the treatment, the proposed model provides a treatment recommendation procedure. Finally, the authors provided a Python library available on the first author's GitHub (Katzman 2017).

### ***Latent-variable Models***

The final model is a latent-variable model based on the Variational Auto-Encoder (VAE) (Kingma 2017; Kingma and Welling 2014) architecture. Louizos et al. (2017) recently suggested a latent variable model for causal inference. The latent variables allow for a more flexible observed variable distribution and intuitively model the hidden patient status. Inspired by this model and by the recommendation of Nazabal et al. (2020), we implemented a latent variable model (Beaulac, Rosenthal, and Hodgson 2018) that adapts the VAE architecture for the purpose of survival analysis. This Survival Analysis Variational Auto-Encoder (SAVAE) uses the latent space to represent the patient true sickness status and can produce individual patient survival function based on their respective covariates which should solve problems (1), (2), and (3).

## **Data Analysis**

### ***Evaluation Metrics***


We will use two different metrics to evaluate the various algorithms, both are well established and they evaluate different properties of the models. First, the concordance index (Harrell, Lee, and Mark 1996) is a metric of accuracy for the ordering of the predicted survival time or hazard. Second, the Brier score (Graf et al. 1999) is a metric similar to the mean squared error but adapted for right-censored observations.

### Concordance Index

The concordance index ( $c$ -index) was proposed by Harrell, Lee, and Mark (1996). It is one of the most popular performance measures for survival problems (Chen et al. 2012; Katzman 2017; Steck et al. 2008) as it elegantly accounts for the censored data. It is defined as the proportion of all usable patient pairs in which the predictions and outcomes are concordant. Pairs are said to be concordant if the predicted event times have a concordant ordering with the observed event times.

Recently Steck et al. used the  $c$ -index directly as part of the optimization procedure (Steck et al. 2008), their paper also elegantly presents the  $c$ -index itself using graphical models as illustrated in Figure 1. In their article, it is defined as the fraction of all pairs of subjects whose predicted survival times are correctly ordered among all subjects that can actually be ordered. We expect a random classification algorithm to achieve a  $c$ -index of 0.5. The further from 0.5 the  $c$ -index is the more concordant pairs of predictions the model has produced. A  $c$ -index of 1 indicates perfect predicted order.

Figure 1 illustrates when we can compute the concordance for a pair of data points; this is represented by an arrow. We can evaluate the order of events if both events are observed. If one of the data points is censored, then concordance can be evaluated if the censoring for the censored point happens after the event for the observed point. If the reverse happens, if both points are



**Figure 1.** Steck et al. (2008) graphical representation of the  $c$ -index computation. Filled circle represents observed points and empty circle represents censored points. This figure illustrates the pairs of points for which an order of events can be established.

censored or if both events happen exactly at the same time then we cannot evaluate the concordance for that pair.

### Brier Score

The Brier score established by Graf et al. (1999) is a performance metric inspired by the mean squared errors (MSE). For a survival model, it is reasonable to try to predict  $P(T > t | X = x) = S(t | X = x)$  the survival probabilities a time  $t$  for a patient with predictors  $x$ . In Graf's notation,  $\hat{\pi}(t | x)$  is the predicted probability of survival at time  $t$  for a patient with characteristics  $x$ . These probabilities are used as predictions of the observed event  $y = \mathbf{1}(T > t)$ . If the data contains no censoring, the simplest definition of the Brier Score would be:

$$BS(t) = \frac{1}{n} \sum_{i=1}^n (\mathbf{1}(T_i > t) - \hat{\pi}(t | x_i))^2 \quad (1)$$

Assuming we have a censoring survival distribution  $G(t) = P(C > t)$  and an associated Kaplan–Meier estimated  $\hat{G}(t)$ . For a given fixed time  $t$  we are facing three different scenarios:

Case 1:  $T_i > t$  and  $\delta_i = 1$  or  $\delta_i = 0$

Case 2:  $T_i < t$  and  $\delta_i = 1$

Case 3:  $T_i < t$  and  $\delta_i = 0$ ,

where  $\delta_i = 1$  if the event is observed and 0 if it is censored. For case 1, the event status is 1 since the patient is known to be alive at time  $t$ ; the resulting contribution to the Brier score is  $(1 - \hat{\pi}(t | x_i))^2$ . For case 2, the event occurred before  $t$  and the event status is equal to  $\mathbf{1}(T_i > t) = 0$  and thus the contribution is  $(0 - \hat{\pi}(t | x_i))^2$ . Finally, for case 3 the censoring occurred before  $t$  and thus the contribution to the Brier score cannot be calculated. To compensate for the loss of information due to censoring, the individual contributions have to be reweighed in a similar way as in the calculation of the Kaplan–Meier estimator leading to the following Brier Score:

$$BS^c(t) = \frac{1}{n} \sum_{i=1}^n ((0 - \hat{\pi}(t | x_i))^2 \mathbf{1}(T_i < t, \delta_i = 1)(1/\hat{G}(T_i)) + (1 - \hat{\pi}(t | x_i))^2 \mathbf{1}(T_i > t)(1/\hat{G}(t))) \quad (2)$$

### Comparative Results

The data set introduced in section 2 was imported in both R (R Core Team 2013) and Python (Van Rossum and Drake 1995). To evaluate the algorithms we randomly divided the data set into 1500 training observations and 212 testing observations. The models were fit using the training observations and the evaluation metrics were computed on the testing observations.

As mentioned in the previous sections, the CoxPH benchmark and the conventional statistical learning models were all tested in the R language (R Core Team 2013). They were relatively easy to use with very little adjustment needed and clear and concise documentation. The computational speed of these algorithms was fast enough on a single CPU so that we could perform 50 trials. The state-of-the-art techniques needed a deeper understanding of the model as they contain many hyper-parameters that require calibration. They were also slower to run on a single CPU.

Figure 2 illustrates Sinaplots (Sidiropoulos et al. 2017) with associated Boxplots of the  $c$ -index for the CoxPH model and the 8 competitors. We used standard boxplots on the background since they are common and easy to understand. The sinaplots superposed on them represent the actual observed metric values and convey information about the distribution of the metrics for a given technique. As mentioned earlier  $c$ -index ranges from 0.5 to 1 where a  $c$ -index of 1 indicates perfect predicted order. According to Figure 2, it seems no model clearly outperforms another. It seems like Random Survival Forests is the best-performing model with relatively small variance and high performance but the difference is not statistically significant.

Since the Brier score is a metric inspired by the mean squared error, it ranges from 0 to 1 and the lower the Brier score is the better the technique. In Figure 3, we once again observe that none of the new techniques significantly





Figure 2. Boxplots and sinaplots of the  $c$ -index (higher the better).





**Figure 3.** Boxplots and sinaplots of Brier scores evaluated at 3 years (lower the better).

outperforms any CoxPH. SAVAE has the lowest Brier score but the difference is not significant when compared to other techniques.

### Takeaways and Recommendations

The previous section demonstrates that the new ML methods offer very little improvement compared to the benchmark CoxPH model according to our two designated performance metrics when patient clinical characteristics that are typically collected in clinical trials are used as predictor variables. This is an important result as we need to evaluate the abilities of ML techniques to solve real-life data problems, and to illuminate the changes in clinical data collection that will have to occur for ML methods to be used to greatest effect in assisting outcome prediction and treatment.

Similar results on real-life data sets are observed in article-presenting methodologies (Fotso 2018; Jinga et al. 2019; Luck et al. 2017) where the proposed techniques provide non-significant improvements over simple models such as CoxPH. Christodoulou et al. (2019) recently performed an exhaustive review of 927 articles that discuss the development of diagnostic or prognostic clinical prediction models for binary outcomes based on clinical data. The authors of the review noted the overall poor comparison methodologies and the lack of significant difference between

a simple logistic regression and state-of-the-art ML techniques in most of the recent years' publications. These results are supported by Hand (2006) who discussed in detail the potential strength of the simple models compared to state-of-the-art ML models. This raises an important question of our case study highlights: is it worth using more complex models for a slight improvement?

The alternatives we proposed in section 3 are all more complicated than CoxPH in various ways. Most of the new techniques require a deeper knowledge of the algorithm behaviors to correctly fix the many hyper-parameters. They can produce less interpretable results due to model complexity, and often require more computing power. Indeed, if the CoxPH model can be fit in seconds, most of the conventional statistical learning models take minutes to fit and the state-of-the-art models take hours. Finally, many of the new techniques are not widely accessible or standardized. As an open language, Python offers very little support to users and the libraries are not maintained, not standardized, and come with dependency issues.

Hand (2006) demonstrates the high relative performances of extremely simple methods compared to complex ones and mathematically justifies his argument. He also discusses how these slight improvements over simple models might be undesirable as they might be attributed to overfitting which would cause reproducibility issues on new data sets. These slight improvements might also be artificial as they were achieved only because the inventors of these techniques were able to obtain through much effort the best performance from their own techniques and not the methods described by others. Overall if the improvements over simple techniques are small, perhaps they are simply not an improvement and this argument seems to be supported by both our case study and the recent review of Christodoulou et al. (2019). We recommend that practitioners keep their expectations low when it comes to some of these new models.

In contrast, significant improvements for diagnostic tasks have been accomplished using AI in recent years (Liu et al. 2017; Rodríguez-Ruiz et al. 2019; Rodríguez-Ruiz et al. 2019) and thus we ask ourselves what caused this difference? There is a major difference in the style of data sets that were available. In the cited articles, images (mammographic, gigapixel pathology image, MRI scans) are analyzed using deep convolutional neural networks (CNN) (Goodfellow, Bengio, and Courville 2016). Models such as CNN were developed because a special type of data was available and none of the current tools was equipped to analyze it. Conventional techniques such as logistic regression or CoxPH are not able to grasp the signal in images, which contains a large number of highly correlated predictors that individually contain close to no information but analyzed together contain a lot. As a matter of fact, the greatest

strength of these models is that they are able to extract a lot of information from a rich, but complicated, data set.

In our case study, the *stratum* predictor was a binary predictor indicating if the patient had a rapid early response to the first rounds of chemotherapy. Computed-tomography (CT) scans of the affected regions were analyzed before and after the first round of treatments and this rich information was transformed into a simple binary variable. This practice is common: even in ongoing trials, patients' characteristics continue to be collected manually (often on paper forms), which dramatically limits the capacity to capture the full range of potentially useful data available for analysis. As new tools are established to extract information from ever growing, both in size and complexity, data sets, clinical trialists have to rethink how they gather data and transform it to make sure that no information is lost in order to utilize these new tools. It seems like extracting and keeping as much information as possible and having a data-centric approach where the model is designed to analyze a specific style of data were some of the factors in the success of CNNs.

## Conclusion

In this article, we have identified a series of statistical and ML techniques that should alleviate some of the flaws of the well-known CoxPH model. These models were tested against a real-life data set and provided little to no improvement according to the *c*-index and the Brier score. Although one might anticipate that these techniques would have increased our prediction abilities, instead the CoxPH performed comparably to modern models. These results are supported by other articles with similar findings.

It would be advantageous to try to theoretically understand when the new techniques should work and when they should not. As it currently stands, authors are not incentivized to discuss the weakness of their techniques and it actually slows scientific progress. It is imperative that we try to understand when some of the newest techniques perform poorly and shed the light on why it is the case. It is also important to understand what made some of these new techniques successful. For example, it seems that CNNs were successful since the model was specifically built for images, a special type of data that was previously hard to handle but contained a large amount of information.

## Disclosure Statement

No potential conflict of interest was reported by the authors.

## Funding

Research reported in this work was supported by the Children's Oncology Group; by the National Cancer Institute of the National Institute of Health under the National Clinical Trials Network (NCTN) Operations Center Grant U10CA180886, the NCTN Statistics and Data Center Grant U10CA180899; the Natural Sciences and Engineering Research Council of Canada (NSERC); and by the Ontario Student Assistance Program (OSAP) (the scholarship provider). The content is solely the responsibility of the authors and does not necessarily represent the official views of the Children's Oncology Group, the National Institute of Health, or St. Baldrick's Foundation.

## References

- Beaulac, C., J. S. Rosenthal, and D. Hodgson. 2018. A deep latent-variable model application to select treatment intensity in survival analysis. *Proceedings of the Machine Learning for Health (ML4H) Workshop at NeurIPS 2018* 2018. [https://arxiv.org/search/?searchtype=report\\_num&query=ML4H/2018/53](https://arxiv.org/search/?searchtype=report_num&query=ML4H/2018/53)
- Bou-Hamad, I., D. Larocque, and H. Ben-Ameur. 2011, January. A review of survival trees. *Statistics Surveys* 5 :44–71. doi:10.1214/09-SS047.
- Breiman, L. 1996. Bagging predictors. *Machine Learning* 24 (2):123–40. doi:10.1007/BF00058655.
- Breiman, L. 2001. Random forests. *Machine Learning* 45 (1):5–32. doi:10.1023/A:1010933404324.
- Breiman, L., J. Friedman, R. Olshen, and C. Stone. 1984. *Classification and regression trees*. Monterey, CA: Wadsworth and Brooks.
- Chen, H.-C., R. L. Kodell, K. F. Cheng, and J. J. Chen. 2012, July 23. Assessment of performance of survival prediction models for cancer prognosis. *BMC Medical Research Methodology* 12 (1):102. doi:10.1186/1471-2288-12-102.
- Christodoulou, E., J. Ma, G. S. Collins, E. W. Steyerberg, J. Y. Verbakel, and B. V. Calster. 2019. A systematic review shows no performance benefit of machine learning over logistic regression for clinical prediction models. *Journal of Clinical Epidemiology* 110:12–22. doi:10.1016/j.jclinepi.2019.02.004.
- Cox, D. R. 1972. Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)* 34 (2):187–220. <http://www.jstor.org/stable/2985181>.
- Cox, D. R. 1975, August. Partial likelihood. *Biometrika* 62 (2):269–76. doi:10.1093/biomet/62.2.269.
- Fotso, S. 2018, January. Deep neural networks for survival analysis based on a multi-task framework. *arXiv E-prints*:arXiv:1801.05512.
- Fotso, S., et al. 2019. PySurvival: Open source package for survival analysis modeling. <https://www.pysurvival.io/>
- Friedman, D. L., L. Chen, S. Wolden, A. Buxton, K. McCarten, T. J. FitzGerald, and C. L. Schwartz. 2014. Dose-intensive response-based chemotherapy and radiation therapy for children and adolescents with newly diagnosed intermediate-risk hodgkin lymphoma: A report from the children's oncology group study ahod0031. *Journal of Clinical Oncology* 32 (32):3651–58. (PMID: 25311218). doi:10.1200/JCO.2013.52.5410.
- Friedman, J., T. Hastie, and R. Tibshirani. 2010. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software* 33 (1):1–22. <http://www.jstatsoft.org/v33/i01/>.

- Goodfellow, I., Y. Bengio, and A. Courville. 2016. *Deep learning*. MIT Press. <http://www.deeplearningbook.org>.
- Graf, E., C. Schmoor, W. Sauerbrei, and M. Schumacher. 1999. Assessment and comparison of prognostic classification schemes for survival data. *Statistics in Medicine* 18 (September):2529–45. doi:10.1002/(SICI)1097-0258(19990915/30)18:17/183.0.CO;2-5.
- Haider, H. 2019. MTLR: Survival prediction with multi-task logistic regression [Computer software manual]. (R package version 0.2.1). <https://CRAN.R-project.org/package=MTLR>.
- Hand, D. J. 2006, February. Classifier technology and the illusion of progress. *Statistical Science* 21 (1):1–14. doi:10.1214/088342306000000060.
- Harrell, F. E., K. L. Lee, and D. B. Mark. 1996. Multivariable prognostic models: Issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Statistics in Medicine* 15 (4):361–87. doi:10.1002/(SICI)1097-0258(19960229)15:4<361::AID-SIM168>3.0.CO;2-4.
- Hastie, T., R. Tibshirani, and J. Friedman. 2009. *The elements of statistical learning*. 2nd ed. New York: Springer.
- Hothorn, T., B. Lausen, A. Benner, and M. Radespiel-Tröger. 2004. Bagging survival trees. *Statistics in Medicine* 23 (1):77–91. doi:10.1002/sim.1593.
- Hothorn, T., K. Hornik, and A. Zeileis. 2006. Unbiased recursive partitioning: A conditional inference framework. *Journal of Computational and Graphical Statistics* 15 (3):651–74. doi:10.1198/106186006X133933.
- Hothorn, T., K. Hornik, C. Strobl, and A. Zeileis. 2019. Party: A laboratory for recursive partytioning [Computer software manual]. (R package version 1.3-3). <https://cran.r-project.org/web/packages/party/index.html>.
- Hothorn, T., P. Bühlmann, S. Dudoit, A. Molinaro, and M. J. Van Der Laan. 2005, December. Survival ensembles. *Biostatistics* 7 (3):355–73. doi:10.1093/biostatistics/kxj011.
- Ishwaran, H., and U. Kogalur. 2019. Fast unified random forests for survival, regression, and classification (rf-src) [Computer software manual]. manual. (R package version 2.9.1). <https://cran.r-project.org/package=randomForestSRC>.
- Ishwaran, H., and U. B. Kogalur. 2010. Consistency of random survival forests. *Statistics & Probability Letters* 80 (13):1056–64. <http://www.sciencedirect.com/science/article/pii/S0167715210000672>.
- Ishwaran, H., U. B. Kogalur, E. H. Blackstone, and M. S. Lauer. 2008, September. Random survival forests. *The Annals of Applied Statistics* 2 (3):841–60. doi:10.1214/08-AOAS169.
- Ishwaran, H., U. B. Kogalur, E. Z. Gorodeski, A. J. Minn, and M. S. Lauer. 2010. High-dimensional variable selection for survival data. *Journal of the American Statistical Association* 105 (489):205–17. doi:10.1198/jasa.2009.tm08622.
- Jinga, B., T. Zhang, Z. Wang, Y. Jina, K. Liua, W. Qiua, . . . C. Lia. 2019. A deep survival analysis method based on ranking. *Artificial Intelligence in Medicine* 98:1–9. doi:10.1016/j.artmed.2019.06.001.
- Katzman, J. (2017). *DeepSurv: Personalized treatment recommender system using a cox proportional hazards deep neural network*. <https://github.com/jaredleekatzman/DeepSurv>
- Katzman, J., U. Shaham, A. Cloninger, J. Bates, T. Jiang, and Y. Kluger. 2018, December. DeepSurv: Personalized treatment recommender system using a cox proportional hazards deep neural network. *BMC Medical Research Methodology*. 18 (1). doi:10.1186/s12874-018-0482-1.
- Kingma, D. P. 2017. Variational inference & deep learning: A new synthesis. Unpublished doctoral dissertation, Universiteit van Amsterdam.
- Kingma D.P. and Welling M. 2014. Auto-Encoding Variational Bayes. In Proceedings of the 2nd International Conference on Learning Representations (ICLR)
- LeBlanc, M., and J. Crowley. 1995. A review of tree-based prognostic models. *Recent Advances in Clinical Trial Design and Analysis* 75:113–24.

- Leblanc, M. E., and J. P. Crowley. 1992. Relative risk trees for censored survival data. *Biometrics* 48 (2):411–25. doi:10.2307/2532300.
- Liu, Y., K. Gadepalli, M. Norouzi, G. E. Dahl, T. Kohlberger, A. Boyko, and M. C. Stumpe. 2017, March. Detecting cancer metastases on gigapixel pathology images. *arXiv E-prints* : arXiv:1703.02442.
- Louizos, C., U. Shalit, J. Mooij, D. Sontag, R. Zemel, and M. Welling. 2017, May. Causal effect inference with deep latent-variable models. . In *Advances in Neural Information Processing Systems* (pp. 6446-6456).
- Luck, M., T. Sylvain, H. Cardinal, A. Lodi, and Y. Bengio. 2017. Deep learning for patient-specific kidney graft survival analysis. *CoRR*, abs/1705.10245. <http://arxiv.org/abs/1705.10245>
- Nazabal, A., Olmos, P. M., Ghahramani, Z., & Valera. (2020). Handling incomplete heterogeneous data using vaes. *Pattern Recognition*, 107501.
- Peters, A., and T. Hothorn. 2019. *ipred: Improved predictors* [Computer software manual]. (R package version 0.9-9). <https://CRAN.R-project.org/package=ipred>.
- R Core Team. 2013. R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. <http://www.R-project.org/>.
- Rodríguez-Ruiz, A., E. Krupinski, -J.-J. Mordang, K. Schilling, S. H. Heywang-Köbrunner, I. Sechopoulos, and R. M. Mann. 2019. Detection of breast cancer with mammography: Effect of an artificial intelligence support system. *Radiology* 290 (2):305–14. (PMID: 30457482). doi:10.1148/radiol.2018181371.
- Rodríguez-Ruiz, A., K. Lång, A. Gubern-Merida, M. Broeders, G. Gennaro, P. Clauser, and I. Sechopoulos. 2019, March. Stand-alone artificial intelligence for breast cancer detection in mammography: Comparison with 101 radiologists. *JNCI: Journal of the National Cancer Institute* 111 (9):916–22. doi:10.1093/jnci/djy222.
- Sidiropoulos, N., S. H. Sohi, N. Rapin, and F. O. Bagger (2017). *sinaplot: An enhanced chart for simple and truthful representation of single observations over multiple classes*. <https://cran.r-project.org/web/packages/sinaplot/vignettes/SinaPlot.html>
- Simon, N., J. Friedman, T. Hastie, and R. Tibshirani. 2011. Regularization paths for cox's proportional hazards model via coordinate descent. *Journal of Statistical Software, Articles* 39 (5):1–13. doi:10.18637/jss.v039.i05.
- Steck, H., Krishnapuram, B., Dehing-Oberije, C., Lambin, P., & Raykar, V. C. (2008). On ranking in survival analysis: Bounds on the concordance index. In *Advances in neural information processing systems* (pp. 1209–1216)
- Strobl, C., A.-L. Boulesteix, A. Zeileis, and T. Hothorn. 2007. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC Bioinformatics* 8 (1):25. doi:10.1186/1471-2105-8-25.
- Therneau, T., B. Atkinson, and B. Ripley (2017). *rpart: Recursive partitioning and regression trees* [Computer software manual]. (R package version 4.1-11). <https://CRAN.R-project.org/package=rpart>
- Van Rossum, G., and F. L. Drake Jr. 1995. *Python tutorial*. Amsterdam, The Netherlands: Centrum voor Wiskunde en Informatica.
- Yu, C.-N., R. Greiner, H.-C. Lin, and V. Baracos. 2011. Learning patient-specific cancer survival distributions as a sequence of dependent regressors. In *Advances in neural information processing systems 24*, ed. J. Shawe Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, 1845–53. Curran Associates, Inc. <http://papers.nips.cc/paper/4210-learning-patient-specific-cancer-survival-distributions-as-a-sequence-of-dependent-regressors.pdf>.
- Zhao, L., and D. Feng. 2019, August. DNNSurv: Deep neural networks for survival analysis using pseudo values. *arXiv E-prints* :arXiv:1908.02337.

## Appendix

**Table A1.** Predictor variables and description.

Variables	Type	Description
agedxyrs	Continuous	Age of the patient at the start of the treatment
gender	Binary	Biological gender
stage	Categorical	Cancer stage ranging from 1 to 4
b_symptoms	Binary	Presence of B symptoms
bulk_disease	Binary	Presence of Bulk disease
extralymphatic_disease	Binary	Presence of Extralymphatic disease
fever	Binary	Presence of recurrent fever
night_sweats	Binary	Presence of night sweats
weight_loss	Binary	Presence of significant weight loss (> 10%)
nodal_aggregate	Binary	Presence of a nodal aggregate
mediastinal_mass	Binary	Presence of a mediastinal mass
esron	Continuous	Erthroctye sedimentation rate (mm/hr)
istnon	Continuous	Number of involved nodal sites
histology	Categorical	Histology (LP,LD,NS,MC, unknown)
albon	Continuous	Albumin (g/dL)
hgbon	Continuous	Hemoglobin(g/dL)
amend	Binary	
stratum	Binary	Rapid early response to first treatment
morpho_icdo	Categorical	ICD-O Morphology codes
RT	Binary	Treatment variable: Radiotherapy
DECA	Binary	Treatment variable: Intensive Chemotherapy