

Neural Fine-Gray: Monotonic neural networks for competing risks

Vincent Jeanselme

*The Alan Turing Institute
MRC Biostatistics Unit, University of Cambridge*

VINCENT.JEANSELME@MRC-BSU.CAM.AC.UK

Chang Ho Yoon

*The Alan Turing Institute
Big Data Institute, University of Oxford*

CHANGHO.YOON@SJC.OX.AC.UK

Brian Tom

Jessica Barrett
MRC Biostatistics Unit, University of Cambridge

BRIAN.TOM@MRC-BSU.CAM.AC.UK

JESSICA.BARRETT@MRC-BSU.CAM.AC.UK

Abstract

Time-to-event modelling, known as survival analysis, differs from standard regression as it addresses *censoring* in patients who do not experience the event of interest. Despite competitive performances in tackling this problem, machine learning methods often ignore other *competing risks* that preclude the event of interest. This practice biases the survival estimation. Extensions to address this challenge often rely on parametric assumptions or numerical estimations leading to sub-optimal survival approximations. This paper leverages constrained monotonic neural networks to model each competing survival distribution. This modelling choice ensures the exact likelihood maximisation at a reduced computational cost by using automatic differentiation. The effectiveness of the solution is demonstrated on one synthetic and three medical datasets. Finally, we discuss the implications of considering competing risks when developing risk scores for medical practice.

Data and Code Availability Experiments are performed on publicly available datasets: Primary Biliary Cholangitis¹ (Therneau et al., 2000), Framingham² (Kannel and McGee, 1979), Synthetic³ (Lee et al., 2018), and the Surveillance, Epidemiology, and End Results Program⁴.

1. Available in the R survival package.
2. Available in the R riskCommunicator package.
3. Available at <https://github.com/chl8856/DeepHit>
4. Available at <https://seer.cancer.gov/>

The code to reproduce the proposed model and the presented results is available on GitHub⁵.

Institutional Review Board (IRB) This research does not require IRB approval as it relies on publicly available datasets from previous studies.

1. Introduction

1.1. Motivation

Survival analysis involves modelling the time to an event of interest, which plays a critical role in medicine to understand disease manifestation, treatment outcomes, and the influence of different risk factors on patient health (Selvin, 2008). This analysis differs from standard regression settings as patients may not experience the outcome of interest over the study period. These *censored* patients inform this regression as they participate in the study event-free until exiting the study. Multiple approaches have been proposed to take advantage of these patients by maximising the likelihood of the observed data.

Often, in medical data, patients may experience events, known as *competing risks*, that preclude the observation of the event of interest. For instance, in modelling the time to cardiac events, patients who die from another condition during the observation period exit the study because of a competing risk. Competing

5. <https://github.com/Jeanselme/NeuralFineGray>

risks remain overlooked despite their prevalence in medicine (Koller et al., 2012; Austin et al., 2016b). Particularly, practitioners frequently consider competing risks as censoring (Austin and Fine, 2017a). This practice breaks the common assumption of non-informative censoring, i.e., censored patients must leave the study for reasons independent of the outcome of interest. Considering competing risks as censoring, therefore, results in misestimating the risk of the event of interest (Fisher and Kanarek, 1974; Leung et al., 1997).

To better tackle the problem of competing risks, one can explicitly model them through the marginal probability of observing each risk, known as the Cumulative Incidence Function (CIF). Estimation of these functions often relies on proportional hazards, parametric assumptions, or numerical integration, potentially resulting in the optimisation of a sub-optimal target misrepresenting the true underlying survival distribution.

1.2. Contribution

This work introduces a novel machine learning model to tackle the problem of competing risks. This approach generalises Rindt et al. (2022) to competing risks, leveraging monotonic neural networks to model cumulative incidence functions. The proposed method tackles the limitations of existing strategies by an exact computation of the likelihood at a lower computational cost.

First, we explore the existing literature before introducing in detail our proposed model. Subsequently, we demonstrate the advantages and limitations of our approach as applied to one synthetic and three real-world medical datasets. Finally, we further investigate the Framingham dataset to underline the importance of considering competing risks in cardiovascular disease risk estimation.

2. Related work

This section summarises the recent progress in machine learning for survival analysis.

2.1. Time-to-event modelling

Survival analysis is an active field of research in the statistical community (Kartsonaki, 2016). Non-parametric (Ishwaran et al., 2008) and parametric (Cox, 2008; Royston, 2001; Cox et al., 2007) models have been introduced to model survival outcomes. Despite these multiple alternatives and considerable proposed extensions, the original Cox proportional-hazards model (Cox, 1972) remains widely used in the medical literature (Stensrud and Hernán, 2020). This semi-parametric approach estimates the impact of covariates on the instantaneous risk of observing an event, i.e., hazard. The model assumes the hazard to take the form of the product of a non-parametric estimate of the population survival and a parametric covariate effect. This assumption is known as proportional hazards and renders tractable the model optimisation for covariate effect estimation.

The machine learning community has extended the Cox model for unknown parametric forms of covariate effect. Specifically, DeepSurv (Katzman et al., 2018) replaces this otherwise parametric component with a neural network. However, this model still assumes proportional hazards that may not hold in real-world medical settings (Stensrud and Hernán, 2020). To relax this assumption, DeepCox (Nagpal et al., 2021c) identifies subgroups using independent Cox models. Each subgroup is characterised by its own non-parametric baseline and covariate effect. At the intersection between DeepCox and parametric models, Nagpal et al. (2021b) model each subgroup with a Weibull distribution parameterised by neural networks to allow end-to-end training. Jeanselme et al. (2022) abandon the parametric and proportional hazards assumption with unconstrained distributions learnt through monotonic networks.

With a focus on predictive performance, DeepHit (Lee et al., 2018) approaches survival as a classification problem where survival prediction time is discretised. The associated task is to predict the interval at which a patient experiences the event. The model’s training procedure consists of a likelihood and a ranking penalty which favours temporally coherent predictions. Extrapolation of this model to infinite time discretisa-

tion resembles an ordinary differential equation (ODE), as proposed in [Danks and Yau \(2022\)](#).

The models above approximate the underlying survival likelihood either through parametric assumptions, discretisation or numerical integration. Recently, [Rindt et al. \(2022\)](#) proposed to overcome this challenge of likelihood estimation by deploying a constrained neural network with a monotonically increasing outcome to obtain the survival function, and, therefore, the exact likelihood. In addition, to show improved performance, the authors demonstrate that one should prefer likelihood optimisation over discriminative performance as the optimal likelihood is obtained for the true underlying survival distribution, i.e., the likelihood is a proper scoring rule. Our study is a generalisation of this work to competing risks, harnessing monotonic neural networks to directly model CIFs.

2.2. Modelling competing risks

Using the aforementioned models without consideration of competing risks would lead to a misestimation of the risk associated with the event of interest ([Schuster et al., 2020](#)). To tackle this issue, one can independently estimate each competing-risk-specific model and combine them to estimate the risk associated with a specific outcome given the non-observation of the other risks, as formulated in the cause-specific Cox model ([Prentice et al., 1978](#)). This independent estimation describes how covariates impact each event risk ([Austin and Fine, 2017b](#)) but may misrepresent the relative effect of these covariates on outcomes ([Austin et al., 2016b](#)) and lead to sub-optimal predictive performance. Alternatively, [Fine and Gray \(1999\)](#) propose to model the sub-hazards, i.e., the probability of observing a given event if the patient has not experienced this event until t , under an assumption of proportionality analogous to the one made in the Cox proportional-hazards model. While providing insights into the link between covariates and risk particularly suitable for prediction ([Austin and Fine, 2017b](#)), this model suffers from two shortcomings: (i) the proportionality assumption impairs its real-world applicability; (ii) this approach can result in an ill-defined survival function ([Austin et al., 2021](#)).

Machine learning approaches have been extended to jointly model competing risks. DeepHit’s time-discretisation results in a straightforward extension in which the output dimension is multiplied by the number of risks ([Lee et al., 2018](#)). Similarly, hierarchical discretisation ([Tjandra et al., 2021](#)) has been proposed. As parametric distributional assumptions result in a closed-form likelihood, [Nagpal et al. \(2021b\)](#) propose to extend their mixture of Weibull distributions and [Bellot and Schaar \(2018\)](#) introduce a Bayesian mixture of Generalised Gamma distributions to tackle competing risk. Under more complex non-parametric likelihoods, numerical integration ([Danks and Yau, 2022](#); [Aastha and Liu, 2020](#)) and pseudo-value approximations ([Rahman et al., 2021](#)) have been proposed. Finally, non-likelihood-based approaches have been introduced such as boosted trees ([Bellot and van der Schaar, 2018](#)) or survival trees ([Schmid and Berger, 2021](#)). However, these methods are optimised towards a Brier-score-like loss.

While survival analysis has received considerable attention in the machine learning community, the problem of competing risks is less well studied ([Wang et al., 2019](#)) and even less applied ([Monterrubio-Gómez et al., 2022](#)), despite being central to medical applications. The existing methodologies to tackle competing risks rely on parametric assumptions, likelihood approximation, or optimise for a score that may misrepresent the true underlying survival distribution. This paper offers a novel competing risk model relying on constrained networks to obtain CIFs as a derivative instead of an integral. This approach results in the exact maximisation of the likelihood by leveraging automatic differentiation.

3. Proposed approach

This section formalises the problem of survival analysis and introduces the proposed model.

3.1. Notation

We model a population of the form $\{x_i, t_i, d_i\}_i$ with x_i the covariates for patient i , $t_i \in \mathbb{R}^+$ the time of end of follow-up and $d_i \in \llbracket 0, R \rrbracket$ its associated cause. If $d_i \in \llbracket 1, R \rrbracket$, the patient left the study due to one of the R considered risks.

Otherwise, the patient is right-censored, i.e., the patient left the study for an *unrelated* reason before any of the events of interest were observed. In this work, we focus on right-censoring, but the model can easily be extended to left-censoring. Note that we assume that experiencing one event precludes the observation of any other.

3.2. Survival quantities

Single risk. In settings with no competing risk, i.e., $R = 1$, one aims to estimate the *survival function* S , the probability of not observing the event of interest before time t , i.e.:

$$S(t|x) := \mathbb{P}(T \geq t|x)$$

Equivalently, one aims to estimate the *cumulative hazard function* $\Lambda(t|x)$ related to S as follows:

$$S(t|x) := \exp[-\Lambda(t|x)] = \exp\left[-\int_0^t \lambda(u|x)du\right]$$

where $\lambda(t|x) = \lim_{\delta t \rightarrow 0} \frac{\mathbb{P}(t < T < t + \delta t, T \geq t, x)}{\delta t}$ is the instantaneous hazard of observing the event of interest, assuming no previous event(s).

Estimating this quantity may rely on maximising the likelihood of the observed data. The assumption of non-informative censoring, i.e., event and censoring times are independent given the covariates, is necessary to express the likelihood. Specifically, each patient i with an observed event contributes to the likelihood, the probability of experiencing the event at t_i without previous events, i.e., $\lambda(t_i|x_i)S(t_i|x_i)$. The likelihood associated with each censored patient is the probability of not experiencing the event until t_i , i.e., $S(t_i|x_i)$. This results in the following log-likelihood:

$$l = \sum_{i, d_i \neq 0} \log \lambda(t_i|x_i) - \sum_i \Lambda(t_i|x_i) \quad (1)$$

Competing risks. In the context of competing risks $R > 1$, a patient may leave a study for reasons correlated with the event of interest. Practitioners often consider these events as censoring and rely on single-risk models. However, this practice breaks the common assumption of non-informative censoring and results in misestimation of the survival function. When other

events may be observed, $S(t|x)$ is defined as the probability of observing none of the competing risks before time t , i.e.:

$$S(t|x) = 1 - \sum_{r \in [1, R]} F_r(t|x)$$

where F_r , the *Cumulative Incidence Function* (CIF) for the event r denotes the probability of observing the event r before time t without prior occurrence of any competing event(s), i.e.:

$$F_r(t|x) = \mathbb{P}(T < t, \text{risk} = r|x) \quad (2)$$

with T , the random variable denoting the time of observation of any event. Note that the CIF can be expressed as an integral of observing the event in an infinitesimal interval given that no other event was observed until t :

$$F_r(t|x) = \int_0^t \lambda_r(u|x) e^{-\int_0^t \sum_r \lambda_r(s)ds} du \quad (3)$$

with $\lambda_r(t|x) = \lim_{\delta t \rightarrow 0} \frac{\mathbb{P}(t < T < t + \delta t, \text{risk} = r | T \geq t, x)}{\delta t}$, the cause-specific hazard, i.e., the instantaneous risk of observing the event r , with no other previous event.

A final quantity of interest is the cause-specific survival $S_r(t|x)$ that expresses the probability of not observing a given outcome r by time t , i.e.,

$$\begin{aligned} S_r(t|x) &= \mathbb{P}((T \geq t) \cup (T < t, \text{risk} \neq r)|x) \\ &= 1 - F_r(t|x) \end{aligned}$$

Similar to the single-risk settings, we maximise the likelihood to estimate F_r . Importantly, we assume non-informative censoring *once controlled* on all identified competing risks. While this assumption is more likely to hold once all competing risks are accounted for, practitioners suspecting its implausibility should perform sensitivity analysis for this assumption (Jackson et al., 2014). Under this assumption, the likelihood can be expressed analogously to (1): patients with an observed event contribute to the likelihood as the probability of observing the event d_i at t_i without observing any events until t_i , i.e., $\lambda_r(t_i|x_i)S(t_i|x_i)$. This quantity is the partial derivative of F_r with respect to t evaluated at t_i . Remaining censored patients influence the likelihood as the probability of observing no event until t_i , i.e., $S(t_i|x_i)$. The competing risks log-likelihood can, therefore, be expressed as:

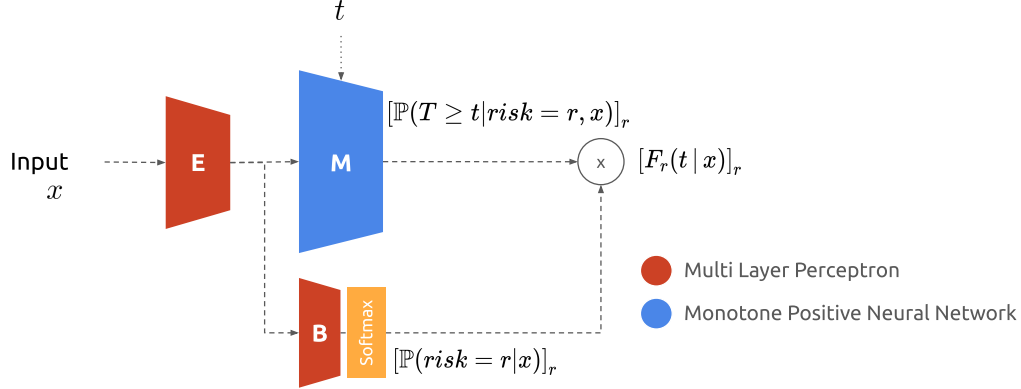


Figure 1: Neural Survival Analysis Architecture. E embeds the covariate(s) x , which are then inputted in the monotonic networks M and balancing network B to estimate the CIFs.

$$l = \sum_{r \in \llbracket 1, R \rrbracket} \sum_{i, d_i=r} \log \left. \frac{\partial F_r(u|x_i)}{\partial u} \right|_{u=t_i} + \sum_{i, d_i=0} \log [1 - \sum_r F_r(t_i|x_i)] \quad (4)$$

One may extend existing models to the competing risks setting by performing the integration in (3). For instance, the cause-specific Cox model (Prentice et al., 1978) consists of Cox models independently trained on each risk, i.e., treating all other outcomes as censored. Then one evaluates the CIF through (3) using the estimated hazards. However, this staged modelling does not jointly consider the outcomes and may misestimate the covariate effects (Van Der Pas et al., 2018). Fine-Gray (Fine and Gray, 1999) overcomes this issue by directly modelling the sub-distribution hazards $h_r(t|x) = \lim_{\delta t \rightarrow 0} \frac{\mathbb{P}(t < T < t + \delta t, \text{risk} = r | (T \geq t) \cup (T < t \cap \text{risk} \neq r), x)}{\delta t}$, relying on a proportionality assumption of these quantities.

Likewise, one can extend machine learning architecture to enable the integration of the CIF and maximise the associated likelihood in (4). However, in the absence of a closed-form expression, this would necessitate numerical integration. This approximation may impact performance with added computational costs for training and predictions. Integration is computationally expensive, whereas derivation can be computed exactly in one backward pass by automatic

differentiation – available in most machine learning libraries. Therefore, our approach reduces the computational cost of the likelihood estimation by modelling F_r and differentiating it to obtain $\lambda_r S$, resulting in the exact computation of all the previously described quantities of interest.

3.3. Architecture

Neural Fine-Gray, illustrated in Figure 1, aims to model $[F_r]_{r \in \llbracket 1, R \rrbracket}$ without relying on numerical integration to tackle the problem of competing risks. We decompose F_r as:

$$F_r(t|x) = \mathbb{P}(\text{risk} = r|x) \cdot \mathbb{P}(T \leq t | \text{risk} = r, x) = B(E(x))_r \cdot [1 - \exp(-t \times M_r(t, E(x)))]$$

Embedding network (E). A first multi-layer perceptron E with inter-layer dropout extracts an embedding \tilde{x} from the covariates x .

Sub-distribution networks ($[M_r]_{r \in \llbracket 1, R \rrbracket}$). The embedding \tilde{x} is inputted in R positive monotonic networks $[M_r]_{r \in \llbracket 1, R \rrbracket}$ representing a lifetime distribution conditioned on one risk r , through the relation $1 - \exp(-t \times M_r(t, \tilde{x})) = \mathbb{P}(T \leq t | x, \text{risk} = r)$. A *positive monotonic neural network* is a network constrained to have its outcome monotonic and positive given its input (see Daniels and Velikova (2010) for theoretical analysis and Lang (2005) for proof of universal approximator). Enforcing these constraints may rely on different transformations of the neural networks’ weights (Omi

et al., 2019; Rindt et al., 2022; Chilinski and Silva, 2020). In our work, we enforce all the neural networks’ weights to be positive through a square function and use a final *SoftPlus* layer to fulfil these constraints. Enforcing positive weights ensures that the outcome increases with the time dimension t . Additionally, enforcing a smooth function ensures a low computational cost and stable optimisation. Note that for model flexibility, we used R monotonic networks. We explore in Appendix B how using one network with R outcomes would impact performance.

Balancing network (B). A multi-layer perceptron B with a final *SoftMax* layer leverages \tilde{x} to balance the probability of observing each risk $B(\tilde{x}) := [\mathbb{P}(\text{risk} = r|x)]_r$. This weighting ensures that the survival function is correctly specified, i.e., $\sum_{r \in [1, R]} F_r(t|x) \leq 1$.

The proposed approach directly models F_r by multiplying the outputs of the distribution and balancing networks. Automatic differentiation of the model’s output results in the derivative $\left. \frac{\partial F_r(u|x_i)}{\partial u} \right|_{u=t_i}$. The model can then be trained end-to-end by maximising the *exact* log-likelihood proposed in Equation (4). By jointly modelling the competing risks, this proposed model is reminiscent of the Fine-Gray approach. The following equation exhibits the link between sub-distribution hazards and CIFs, i.e., between the standard and neural Fine-Gray models:

$$h_r(t|x) = \frac{1}{1 - F_r(t|x)} \cdot \left. \frac{\partial F_r(u|x)}{\partial u} \right|_{u=t}$$

Remark 1 *Shchur et al. (2020) raise a limitation of monotonic neural networks that may attribute non-null density to negative times, i.e., $F_r(t = 0|x) \neq 0$. In contrast to Omi et al. (2019); Rindt et al. (2022), we model $\mathbb{P}(T \leq t|\text{risk} = r, x)$ as $1 - \exp(-t \times M_r(t, \tilde{x}))$ instead of $M_r(t, \tilde{x})$ to address this issue.*

Remark 2 *The proposed methodology is a generalisation of the survival model Sumo-Net (Rindt et al., 2022) that estimates S in the single-risk setting. If $R = 1$, then $F_r = 1 - S$ and $B_r = \mathbb{1}$. In this context, the proposed approach results in Sumo-Net. Moreover, the architecture*

resembles the one proposed in DeSurv (Danks and Yau, 2022) while avoiding numerical integration.

3.4. Computational complexity

Our modelling choices result in the exact computation of the likelihood. However, the other methodologies relying on integral approximation and outcome discretisation converge towards F_r in the upper limit, i.e., when increasing the number of point estimates, or using a finer discretisation. One may therefore question the advantage of the proposed methodology. In this section, we compare the complexity in estimating the CIF and likelihood for DeSurv (Danks and Yau, 2022), the closest method to our proposed model, and NeuralFG.

DeSurv (Danks and Yau, 2022). This approach models $F_r(t|x)$ as $\text{Tanh}(v(x, t))$ with v being the solution to the ODE defined as $\left. \frac{\partial v(x, u)}{\partial u} \right|_{u=t} = g(x, t)$ and $v(x, 0) = 0$ with g , a neural network. For efficiency, the authors propose a Gauss-Legendre quadrature to solve the ODE and obtain v . This approximation necessitates n evaluations of g at defined times $[t_j(t)]_{j \in [1, n]}$ weighted by the associated $[w_j]_{j \in [1, n]}$ (see Press et al. (2007) for a detailed description of Gauss-Legendre quadrature). Each forward pass estimates $\left. \frac{\partial v(x, u)}{\partial u} \right|_{u=t_j(t)}$ at the points used to approximate the integral, then

$$\hat{F}_r(t|x) = \text{Tanh} \left(\frac{t}{2} \sum_{j \in [1, n]} w_j g(x, \frac{t}{2} t_j(t)) \right)$$

DeSurv’s computational cost. Computation of F_r relies on n forward passes through the network. Moreover, the estimation of $\left. \frac{\partial \hat{F}_r(u|x_i)}{\partial u} \right|_{u=t_i}$ necessary to compute the competing risk likelihood is $g(x, t_i)(1 - \text{Tanh}(\hat{F}_r(t_i|x)))^2$, i.e., $n + 1$ forward passes. The likelihood has a $\mathcal{O}(nN)$ computational complexity with N the number of patients in the study.

NeuralFG’s computational cost. F_r is estimated in one forward pass and $\left. \frac{\partial \hat{F}_r(u|x_i)}{\partial u} \right|_{u=t_i}$ in

Dataset	Observations	Features	Primary	Competing risk	Censored
PBC	312	25	Death (44.87 %)	Transplant (9.29 %)	45.83 %
Framingham	4,434	18	CVD (26.09 %)	Death (17.75 %)	56.16 %
Synthetic	30,000	12	* (25.33 %)	* (24.67 %)	50.00 %
SEER	658,354	23	BC (16.51 %)	CVD (5.69 %)	77.80 %

Table 1: Datasets characteristics

one backward pass. Assuming the same computational cost for forward and backward passes, the likelihood estimation has a $\mathcal{O}(2N)$ complexity. Our proposed methodology, therefore, presents more than an $n/2$ computational gain compared to DeSurv in estimating the likelihood used for training, and an n gain in inferring F_r .

4. Experiments

This section introduces the datasets and experimental settings.

4.1. Datasets

We explore the model performance on four datasets with competing risks:

- PBC (Therneau et al., 2000) comprises 25 covariates in 312 patients over a 10-year randomised control trial to measure the impact of D-penicillamine on Primary Biliary Cholangitis (PBC). Death on the waiting list is the primary outcome with transplant being a competing risk.
- Framingham (Kannel and McGee, 1979) is a cohort study gathering 18 longitudinal measurements over 20 years. Our analysis focuses on the first observed covariates of 4,434 patients to model cardiovascular disease (CVD) risk. Death from other causes is treated as a competing risk.
- Synthetic (Lee et al., 2018), this dataset consists of 30,000 synthetic patients with 12 covariates following exponential event time distributions, non-linearly dependent on the covariates.
- SEER⁶: the Surveillance, Epidemiology, and End Results Program gathers covariates and

outcomes of patients diagnosed with breast cancer between 1992 and 2017. Following the preprocessing proposed by Lee et al. (2018); Danks and Yau (2022), we select 658,354 patients and 23 covariates describing the patient demographics and disease characteristics at diagnosis. Death from breast cancer (BC) is our primary outcome, with CVD, a competing risk.

Table 1 summarises the datasets’ characteristics with the respective proportion of outcome and censoring.

4.2. Baseline models

The proposed Neural Fine-Gray (**NeuralFG**) was compared against six strategies. First, we considered the well-established cause-specific Cox model (**CS Cox** Prentice et al. (1978)) and **Fine-Gray** model (Fine and Gray, 1999) with a linear parametric form for the covariate effect. The cause-specific Cox model models each cause independently using a Cox proportional-hazards model, while Fine-Gray models the sub-hazard functions assuming proportional sub-hazards.

Thereafter, we compare state-of-the-art competing risk survival neural networks proposed in the machine learning literature. First, Deep Survival Machine (**DSM**, Nagpal et al. (2021b)) consists of a mixture of Weibull distributions parameterised by neural networks. Each point is then assigned to these distributions through an assignment network. Using parametric distributions results in a closed-form likelihood in the competing risks setting. **DeepHit** (Lee et al., 2018) discretises the survival horizon and leverages a multi-head network to associate each patient to the interval corresponding to its observed event time and type. Each head of the network is associated with one cause as in the proposed NeuralFG. The time-discretisation results in a discrete likelihood further penalised by a C-index-like regu-

6. <https://seer.cancer.gov/>

	Risk	Model	C-Index (<i>Larger is better</i>)			Brier Score (<i>Smaller is better</i>)		
			$q_{0.25}$	$q_{0.50}$	$q_{0.75}$	$q_{0.25}$	$q_{0.50}$	$q_{0.75}$
PBC	Death	NeuralFG	0.810 (0.079)	0.795 (0.114)	0.762 (0.123)	0.099 (0.028)	0.140 (0.020)	0.169 (0.050)
		DeepHit	0.822 (0.099)	0.844 (0.036)	0.782 (0.033)	<i>0.090</i> (0.030)	0.132 (0.013)	0.180 (0.021)
		DeSurv	0.821 (0.089)	0.837 (0.050)	0.815 (0.068)	0.088 (0.022)	0.113 (0.011)	0.136 (0.047)
		DSM	0.867 (0.065)	0.864 (0.037)	0.828 (0.052)	0.091 (0.039)	0.124 (0.015)	0.161 (0.022)
		Fine-Gray	0.831 (0.136)	<i>0.852</i> (0.045)	<i>0.816</i> (0.059)	0.091 (0.042)	<i>0.103</i> (0.009)	0.150 (0.038)
		CS Cox	<i>0.833</i> (0.125)	0.851 (0.040)	0.811 (0.065)	0.091 (0.038)	0.102 (0.008)	<i>0.148</i> (0.038)
Framingham	CVD	NeuralFG	0.872 (0.024)	0.812 (0.029)	0.782 (0.018)	<i>0.050</i> (0.003)	0.095 (0.010)	0.128 (0.004)
		DeepHit	0.855 (0.026)	0.781 (0.026)	0.743 (0.014)	0.053 (0.003)	0.102 (0.007)	0.141 (0.002)
		DeSurv	0.872 (0.027)	<i>0.807</i> (0.031)	0.775 (0.022)	0.049 (0.005)	0.095 (0.009)	<i>0.129</i> (0.003)
		DSM	<i>0.866</i> (0.023)	0.806 (0.023)	<i>0.778</i> (0.014)	0.057 (0.005)	0.104 (0.006)	0.141 (0.002)
		Fine-Gray	0.842 (0.025)	0.794 (0.024)	0.772 (0.015)	0.057 (0.006)	0.099 (0.007)	0.131 (0.003)
		CS Cox	0.845 (0.020)	0.798 (0.022)	0.774 (0.015)	0.056 (0.006)	<i>0.098</i> (0.007)	0.131 (0.003)
Synthetic	1	NeuralFG	<i>0.791</i> (0.013)	<i>0.754</i> (0.013)	0.715 (0.011)	0.068 (0.003)	<i>0.125</i> (0.004)	0.192 (0.005)
		DeepHit	0.783 (0.012)	0.747 (0.013)	<i>0.714</i> (0.008)	0.079 (0.003)	0.136 (0.002)	<i>0.212</i> (0.003)
		DeSurv	0.793 (0.013)	0.756 (0.014)	<i>0.714</i> (0.014)	0.068 (0.002)	0.124 (0.004)	0.192 (0.004)
		DSM	0.776 (0.013)	0.742 (0.013)	0.710 (0.013)	<i>0.073</i> (0.002)	0.139 (0.002)	0.220 (0.003)
		Fine-Gray	0.611 (0.014)	0.587 (0.007)	0.568 (0.009)	0.078 (0.002)	0.159 (0.003)	0.241 (0.002)
		CS Cox	0.609 (0.015)	0.586 (0.006)	0.568 (0.009)	0.078 (0.002)	0.159 (0.003)	0.240 (0.002)
SEER	BC	NeuralFG	<i>0.893</i> (0.002)	<i>0.855</i> (0.001)	<i>0.815</i> (0.001)	0.038 (0.000)	0.069 (0.001)	0.101 (0.000)
		DeepHit	0.899 (0.002)	0.860 (0.001)	0.818 (0.001)	0.038 (0.000)	<i>0.070</i> (0.000)	<i>0.102</i> (0.001)
		DeSurv	0.892 (0.003)	0.852 (0.002)	0.813 (0.001)	0.038 (0.000)	<i>0.070</i> (0.000)	<i>0.102</i> (0.001)
		DSM	0.884 (0.001)	0.842 (0.002)	0.805 (0.002)	<i>0.039</i> (0.000)	0.076 (0.001)	0.112 (0.000)
		Fine-Gray	0.836 (0.003)	0.786 (0.003)	0.742 (0.002)	0.043 (0.001)	0.081 (0.000)	0.118 (0.000)
		CS Cox	0.837 (0.003)	0.786 (0.003)	0.742 (0.002)	0.042 (0.001)	0.081 (0.000)	0.118 (0.000)

Table 2: Comparison of model performance by means (standard deviations) across 5-fold cross-validation. Best performances are in **bold**, second best in *italics*. *NeuralFG* is the model introduced in this paper.

larisation for model training. Closer to our work, **DeSurv** (Danks and Yau, 2022) approaches F_r as the solution to an ODE.

4.3. Experimental settings

The analysis relies on 5-fold cross-validation with 10% of each training set left aside for hyperparameter tuning. Random search is used on the following grid over 100 iterations: learning rate (10^{-3} or 10^{-4}), batch size (100, 250, except for SEER: 1, 000 or 5, 000), dropout rate (0, 0.25, 0.5 or 0.75), number of layers ($\llbracket 1, 4 \rrbracket$) and nodes (25 or 50). All activation functions are fixed to *Tanh* to ensure a properly defined derivative – note that any \mathcal{C}^1 activation would work. All models are optimised using an Adam optimiser (Kingma

and Ba, 2015) over 1,000 epochs, with an early stopping criterion computed on a 10% left-aside subset of the training set.

Other methods are optimised over the same grid (if applicable). Additionally, we explore both Log-Normal and Weibull distributions for DSM and use 10,000 warm-up iterations to estimate the parametric form closest to the average survival as proposed in the original paper (Nagpal et al., 2021b). For DeSurv, we followed the original paper’s recommendation of a 15-degree Gauss-Legendre quadrature to estimate the CIFs. In Appendix C.1, we further investigate how increasing the number of point estimates impacts performance. We use a similar approximation for DeepHit with a 15-split time discretisation. Finally, for a fair compari-

Death	Model	C-Index (<i>Larger is better</i>)			Brier Score (<i>Smaller is better</i>)		
		$q_{0.25}$	$q_{0.50}$	$q_{0.75}$	$q_{0.25}$	$q_{0.50}$	$q_{0.75}$
CVD	Competing	0.872 (0.024)	0.812 (0.029)	0.782 (0.018)	0.050 (0.003)	0.095 (0.010)	0.128 (0.004)
	Non-Competing	0.862 (0.029)	0.807 (0.032)	0.780 (0.020)	0.053 (0.004)	0.099 (0.011)	0.129 (0.005)
Death	Competing	0.745 (0.055)	0.717 (0.038)	0.713 (0.022)	0.027 (0.003)	0.070 (0.004)	0.112 (0.005)
	Non-Competing	0.741 (0.053)	0.718 (0.045)	0.719 (0.025)	0.027 (0.003)	0.071 (0.002)	0.109 (0.004)

Table 3: Modelling competing risk - means (standard deviations) across the 5-fold cross-validation.

son, we double the number of possible layers for architectures without embedding networks.

4.4. Evaluation metrics

As per current practice in survival literature, we used the time-dependent Brier score (Graf et al., 1999) to quantify calibration, and the C-index (Antolini et al., 2005) for discrimination at the dataset-specific 0.25, 0.5 and 0.75 quantiles of the uncensored population event times (See Appendix A.1 for data characteristics, A.2 for further description of the metrics and A.4 for the cumulative version of these metrics). Means and standard deviations are computed over the 5 folds of cross-validation.

5. Results

Table 2 summarises the calibration and discriminative performance of the analysed models on the primary outcome (see Appendix A.3 for the performances on the competing risk).

5.1. Model’s strengths

NeuralFG demonstrates lower or equal Brier scores than other state-of-the-art machine learning models across the majority of datasets and time horizons.

While DSM presents good discriminative performances, this edge is not reflected in its calibration. This observation indicates that parametric assumptions may result in estimated survival functions discriminative of the outcome but further from the underlying survival distribution. Deep-Hit penalisation results in better C-Index values but hurts model calibration, with misaligned discrimination and calibration throughout the different datasets.

Finally, performances are comparable to DeSurv. However, DeSurv’s likelihood approximation multiplies its computational cost by the numerical integration complexity (see Appendix C.2 for a comparison of training speed on the Framingham dataset). NeuralFG, therefore, achieves state-of-the-art performance while avoiding computationally-expensive approximations.

5.2. Model’s limitations

The proposed methodology has lower performance on the PBC dataset, which notably comprises a limited amount of data. In small-data settings, practitioners should prefer simpler models to avoid overfitting.

For instance, the linear Fine-Gray and CS Cox models result in competitive performances on PBC. However, this linearity assumption hurts performance under more complex covariate effects as in the SEER and Synthetic datasets. Note that leveraging domain expertise could enhance performance through the addition of interactions and the use of alternative models. However, these approaches deviate from the automated discovery of interactions facilitated by neural networks. Similarly, the parametric assumption of DSM results in the best discrimination in PBC, but it under-performs under more complex survival distributions.

Furthermore, the DeSurv model performs better than the proposed methodology on PBC. This may reflect that approximating the likelihood can regularise model training, which is beneficial in the context of small data.

5.3. Modelling vs ignoring competing risks

This last section explores the importance of modelling competing risks in the Framingham dataset. First, we present the performance differences between the proposed model in comparison to the same architecture maximising the cause-specific likelihoods. Then, we explore which subgroups of the population most benefit from this modelling. Finally, we study how guidelines would differ under the proposed NeuralFG and its non-competing alternative.

Why account for competing risks? To measure how modelling competing risks impacts performance, while ensuring the *same number of parameters*, we propose to use the same architecture presented in Section 3.3 whilst maximising the sum of the cause-specific likelihoods, i.e.:

$$l = \sum_r \left[\sum_{i, d_i=r} \log \lambda_r(t_i | \tilde{x}_i) - \sum_i \Lambda_r(t_i | \tilde{x}_i) \right]$$

Each monotonic network, therefore, models the cumulative hazard function for risk r , Λ_r , by maximising the likelihood of one cause whilst considering the rest of the population as censored, relying on a shared embedding \tilde{x} . Automatic differentiation outputs $[\lambda_r]_{r \in [1, R]}$. Table 3 summarises the discrimination and calibration differences in the non-competing survival $e^{-\Lambda_r(t|x)}$ obtained with this model and the previously described NFG’s cause-specific survival $1 - F_r(t|x)$. Note how modelling competing risks significantly improves performance for the primary outcome of interest, CVD, without significant differences for the competing risk. Since patients who die from other causes during the study period do not present the same risk of CVD as patients remaining in the study, not accounting for all-cause mortality results in a misestimation of CVD risk.

Who may benefit? One can explore which subgroups benefit the most from modelling competing risks. Intuitively, patients who are the most likely to suffer from competing risks may benefit the most from this modelling. Table 4 illustrates this with older patients benefiting the most from modelling death as a competing risk.

Age	Brier Score Difference		
	$q_{0.25}$	$q_{0.50}$	$q_{0.75}$
< 40	-0.000 (0.000)	-0.001 (0.002)	0.000 (0.005)
40-50	-0.001 (0.001)	-0.002 (0.003)	-0.002 (0.001)
50-60	-0.003 (0.005)	-0.004 (0.003)	-0.006 (0.007)
60+	-0.013 (0.011)	-0.022 (0.018)	-0.007 (0.024)

Table 4: Calibration differences - Means and standard deviations over 5-fold cross-validation. *Larger negative values correspond to better calibration for the competing risk model.*

What is the impact on medical practice?

The Framingham dataset was used to model the eponymous 10-year cardiovascular disease (CVD) risk score (Wilson et al., 1998). This score guides clinical practice in preventatively treating patients, usually with a combination of cholesterol-lowering therapy, e.g., statins, and holistic treatment of other CVD risk factors (Bosomworth, 2011). To minimise overtreatment and adverse side effects, accurate risk estimates are critical for targeting the population most at risk so as to maximise the benefit-risk ratio (Mangione et al., 2022). However, the original Framingham score relies on a non-competing risk model (Mangione et al., 2022; van Kempen et al., 2014).

Clinical treatment often relies on a discretisation of this risk (Bosomworth, 2011): low, intermediate and high risk, at < 10%, 10 – 20% and > 20% chance, respectively, of observing a CVD event in the following 10 years. Current guidelines in the United States suggest placing all patients with $\geq 10\%$ risk on cholesterol-lowering drugs (Mangione et al., 2022). Furthermore, in the US alone, several million patients are on these medications (Wall et al., 2018). Therefore, even modest shifts in patient risk classification could, at scale, amount to considerable numbers either inappropriately receiving preventative treatment or inappropriately receiving none. To demonstrate how considering competing risks can fundamentally alter such risk profiling, we present in Table 5 the reclassification matrices of risk levels given competing and non-competing NeuralFG differentiated by observed outcomes for patients aged 50 or over. For instance, note that 251 deemed intermediate-to-

Table 5: Reclassification matrices between competing and non-competing risk scores for patients older than 50. *Red (resp. blue)* shows when the competing risks score is less aligned with the 10-year observed outcome than the non-competing model (resp. more aligned). Note that censored patients are ignored.

		Non - Competing			Total
Risk		Low	Inter.	High	
Comp.	Low	502	228	23	753
	Intermediate	2	189	229	420
	High	0	9	314	323
	Total	504	426	566	1496

		Non - Competing			Total
Risk		Low	Inter.	High	
Comp.	Low	23	28	5	56
	Intermediate	1	37	41	79
	High	2	4	248	254
	Total	26	69	294	389

high risk by the non-competing risks model are reclassified as lower risk by the competing-risks model, who, in turn, could have avoided the initiation of therapy. These results echo the medical literature’s findings of risk misestimation due to the non-consideration of competing risks in this risk score (Lloyd-Jones et al., 2004; van Kempen et al., 2014). More accurate simulations to estimate the potential lives saved and harmed through such reclassification is beyond the scope of this article but could provide insight into the possible consequences of considering competing risks. In summary, using a non-competing risk score would have important clinical consequences of over- and under-treatment (Schuster et al., 2020). More predictive models accounting for competing risks must be preferred to ensure better care.

6. Conclusion

This work provides a solution to address competing risks that preclude the observation of the outcome of interest, often present in medical applications. We introduce Neural Fine-Gray, a monotonic neural network architecture, to tackle the problem of competing risks in survival modelling. The model outputs the cumulative incidence functions and, consequently, allows the exact likelihood computation. Importantly, this architecture choice achieves competitive performance while avoiding the parametric assumptions or computationally expensive approximations made by state-of-the-art survival neural

(b) Patients with an observed event during the 10-year follow-up.

networks. Further analysis of the Framingham dataset contributes to the literature, inviting practitioners to use competing-risk modelling in risk score development for improved care (Abdel-Qadir et al., 2018; Austin et al., 2016a; Lloyd-Jones et al., 2004; Schuster et al., 2020).

Our future work will (i) extend this architecture to model other modalities such as time series as in Nagpal et al. (2021a) and, (ii) explore medically interpretable survival clusters as presented in Jeanselme et al. (2022); Nagpal et al. (2022a).

Acknowledgments

This work was supported by The Alan Turing Institute’s Enrichment Scheme and partially funded by UKRI Medical Research Council (MC_UU_00002/5 and MC_UU_00002/2).

References

- Pengyu Huang Aastha and Yan Liu. Deepcompete: A deep learning approach to competing risks in continuous time domain. In *AMIA Annual Symposium Proceedings*, volume 2020, page 177. American Medical Informatics Association, 2020.
- Husam Abdel-Qadir, Jiming Fang, Douglas S Lee, Jack V Tu, Eitan Amir, Peter C Austin, and Geoffrey M Anderson. Importance of considering competing risks in time-to-event analyses: application to stroke risk in a retrospective cohort study of elderly patients with atrial

- fibrillation. *Circulation: Cardiovascular Quality and Outcomes*, 11(7):e004580, 2018.
- Laura Antolini, Patrizia Boracchi, and Elia Biganzoli. A time-dependent discrimination index for survival data. *Statistics in medicine*, 24(24):3927–3944, 2005.
- Peter C Austin and Jason P Fine. Accounting for competing risks in randomized controlled trials: a review and recommendations for improvement. *Statistics in medicine*, 36(8):1203–1209, 2017a.
- Peter C Austin and Jason P Fine. Practical recommendations for reporting fine-gray model analyses for competing risk data. *Statistics in medicine*, 36(27):4391–4400, 2017b.
- Peter C Austin, Douglas S Lee, Ralph B D’Agostino, and Jason P Fine. Developing points-based risk-scoring systems in the presence of competing risks. *Statistics in medicine*, 35(22):4056–4072, 2016a.
- Peter C Austin, Douglas S Lee, and Jason P Fine. Introduction to the analysis of survival data in the presence of competing risks. *Circulation*, 133(6):601–609, 2016b.
- Peter C Austin, Ewout W Steyerberg, and Hein Putter. Fine-gray subdistribution hazard models to simultaneously estimate the absolute risk of different event types: cumulative total failure probability may exceed 1. *Statistics in Medicine*, 40(19):4200–4212, 2021.
- Alexis Bellot and Mihaela Schaar. Tree-based bayesian mixture model for competing risks. In *International Conference on Artificial Intelligence and Statistics*, pages 910–918. PMLR, 2018.
- Alexis Bellot and Mihaela van der Schaar. Multitask boosting for survival analysis with competing risks. *Advances in Neural Information Processing Systems*, 31, 2018.
- N John Bosomworth. Practical use of the framingham risk score in primary prevention: Canadian perspective. *Canadian Family Physician*, 57(4):417–423, 2011.
- Pawel Chilinski and Ricardo Silva. Neural likelihoods via cumulative distribution functions. In *Conference on Uncertainty in Artificial Intelligence*, pages 420–429. PMLR, 2020.
- Christopher Cox. The generalized f distribution: an umbrella for parametric survival analysis. *Statistics in medicine*, 27(21):4301–4312, 2008.
- Christopher Cox, Haitao Chu, Michael F Schneider, and Alvaro Munoz. Parametric survival analysis and taxonomy of hazard functions for the generalized gamma distribution. *Statistics in medicine*, 26(23):4352–4374, 2007.
- David R Cox. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2):187–202, 1972.
- Hennie Daniels and Marina Velikova. Monotone and partially monotone neural networks. *IEEE Transactions on Neural Networks*, 21(6):906–917, 2010.
- Dominic Danks and Christopher Yau. Derivative-based neural modelling of cumulative distribution functions for survival analysis. In *International Conference on Artificial Intelligence and Statistics*, pages 7240–7256. PMLR, 2022.
- Jason P Fine and Robert J Gray. A proportional hazards model for the subdistribution of a competing risk. *Journal of the American statistical association*, 94(446):496–509, 1999.
- Lloyd Fisher and Paula Kanarek. Presenting censored survival data when censoring and survival times may not be independent. *Reliability and Biometry*, pages 303–326, 1974.
- Erika Graf, Claudia Schmoor, Willi Sauerbrei, and Martin Schumacher. Assessment and comparison of prognostic classification schemes for survival data. *Statistics in medicine*, 18(17-18):2529–2545, 1999.
- Hemant Ishwaran, Udaya B Kogalur, Eugene H Blackstone, and Michael S Lauer. Random survival forests. 2008.
- Dan Jackson, Ian R White, Shaun Seaman, Hannah Evans, Kathy Baisley, and James Carpenter. Relaxing the independent censoring assumption in the cox proportional hazards model using multiple imputation. *Statistics in medicine*, 33(27):4681–4694, 2014.

- Vincent Jeanselme, Brian Tom, and Jessica Barrett. Neural survival clustering: Non-parametric mixture of neural networks for survival clustering. In *Conference on Health, Inference, and Learning*, pages 92–102. PMLR, 2022.
- William B Kannel and Daniel L McGee. Diabetes and cardiovascular disease: the framingham study. *Jama*, 241(19):2035–2038, 1979.
- Christiana Kartsonaki. Survival analysis. *Diagnostic Histopathology*, 22(7):263–270, 2016.
- Jared L Katzman, Uri Shaham, Alexander Cloninger, Jonathan Bates, Tingting Jiang, and Yuval Kluger. Deepsurv: personalized treatment recommender system using a cox proportional hazards deep neural network. *BMC medical research methodology*, 18(1):1–12, 2018.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015*, 2015.
- Michael T Koller, Heike Raatz, Ewout W Steyerberg, and Marcel Wolbers. Competing risks and the clinical community: irrelevance or ignorance? *Statistics in medicine*, 31(11-12): 1089–1097, 2012.
- Bernhard Lang. Monotonic multi-layer perceptron networks as universal approximators. In *Artificial Neural Networks: Formal Models and Their Applications–ICANN 2005: 15th International Conference, Warsaw, Poland, September 11-15, 2005. Proceedings, Part II 15*, pages 31–37. Springer, 2005.
- Changhee Lee, William Zame, Jinsung Yoon, and Mihaela van der Schaar. Deephit: A deep learning approach to survival analysis with competing risks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Kwan-Moon Leung, Robert M Elashoff, and Abdelmonem A Afifi. Censoring issues in survival analysis. *Annual review of public health*, 18(1): 83–104, 1997.
- Donald M Lloyd-Jones, Peter WF Wilson, Martin G Larson, Alexa Beiser, Eric P Leip, Ralph B D’Agostino, and Daniel Levy. Framingham risk score and prediction of lifetime risk for coronary heart disease. *The American journal of cardiology*, 94(1):20–24, 2004.
- Carol M Mangione, Michael J Barry, Wanda K Nicholson, Michael Cabana, David Chelmsow, Tumaini Rucker Coker, Esa M Davis, Katrina E Donahue, Carlos Roberto Jaén, Martha Kubik, et al. Statin use for the primary prevention of cardiovascular disease in adults: Us preventive services task force recommendation statement. *JAMA*, 328(8):746–753, 2022.
- Karla Monterrubio-Gómez, Nathan Constantine-Cooke, and Catalina A Vallejos. A review on competing risks methods for survival analysis. *arXiv preprint arXiv:2212.05157*, 2022.
- Chirag Nagpal, Vincent Jeanselme, and Artur Dubrawski. Deep parametric time-to-event regression with time-varying covariates. In Russell Greiner, Neeraj Kumar, Thomas Alexander Gerds, and Mihaela van der Schaar, editors, *Proceedings of AAAI Spring Symposium on Survival Prediction - Algorithms, Challenges, and Applications 2021*, volume 146 of *Proceedings of Machine Learning Research*, pages 184–193. PMLR, 22–24 Mar 2021a.
- Chirag Nagpal, Xinyu Li, and Artur Dubrawski. Deep survival machines: Fully parametric survival regression and representation learning for censored data with competing risks. *IEEE Journal of Biomedical and Health Informatics*, 2021b.
- Chirag Nagpal, Steve Yadlowsky, Negar Roshtamzadeh, and Katherine Heller. Deep cox mixtures for survival regression. *Machine Learning for Healthcare Conference*, 2021c.
- Chirag Nagpal, Mononito Goswami, Keith Dufendach, and Artur Dubrawski. Counterfactual phenotyping with censored time-to-events. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD ’22, page 3634–3644, 2022a.
- Chirag Nagpal, Willa Potosnak, and Artur Dubrawski. auton-survival: an open-source

- package for regression, counterfactual estimation, evaluation and phenotyping with censored time-to-event data. In Zachary Lipton, Rajesh Ranganath, Mark Sendak, Michael Sjöding, and Serena Yeung, editors, *Proceedings of the 7th Machine Learning for Healthcare Conference*, volume 182 of *Proceedings of Machine Learning Research*, pages 585–608. PMLR, 05–06 Aug 2022b.
- Takahiro Omi, Kazuyuki Aihara, et al. Fully neural network based model for general temporal point processes. In *Advances in Neural Information Processing Systems*, pages 2122–2132, 2019.
- Sebastian Pölsterl. scikit-survival: A library for time-to-event analysis built on top of scikit-learn. *The Journal of Machine Learning Research*, 21(1):8747–8752, 2020.
- Ross L Prentice, John D Kalbfleisch, Arthur V Peterson Jr, Nancy Flourney, Vernon T Farewell, and Norman E Breslow. The analysis of failure times in the presence of competing risks. *Biometrics*, pages 541–554, 1978.
- William H Press, Saul A Teukolsky, William T Vetterling, and Brian P Flannery. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press, 2007.
- Md Mahmudur Rahman, Koji Matsuo, Shinya Matsuzaki, and Sanjay Purushotham. Deeppseudo: pseudo value based deep learning models for competing risk analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 479–487, 2021.
- David Rindt, Robert Hu, David Steinsaltz, and Dino Sejdinovic. Survival regression with proper scoring rules and monotonic neural networks. In *International Conference on Artificial Intelligence and Statistics*, pages 1190–1205. PMLR, 2022.
- Patrick Royston. Flexible parametric alternatives to the cox model, and more. *The Stata Journal*, 1(1):1–28, 2001.
- Matthias Schmid and Moritz Berger. Competing risks analysis for discrete time-to-event data. *Wiley Interdisciplinary Reviews: Computational Statistics*, 13(5):e1529, 2021.
- Noah A Schuster, Emiel O Hoogendijk, Almar AL Kok, Jos WR Twisk, and Martijn W Heymans. Ignoring competing events in the analysis of survival data may lead to biased results: a nonmathematical illustration of competing risk analysis. *Journal of clinical epidemiology*, 122:42–48, 2020.
- Steve Selvin. *Survival analysis for epidemiologic and medical research*. Cambridge University Press, 2008.
- Oleksandr Shchur, Marin Biloš, and Stephan Günnemann. Intensity-free learning of temporal point processes. In *International Conference on Learning Representations*, 2020.
- Mats J Stensrud and Miguel A Hernán. Why test for proportional hazards? *Jama*, 323(14):1401–1402, 2020.
- Terry M Therneau, Patricia M Grambsch, Terry M Therneau, and Patricia M Grambsch. *The cox model*. Springer, 2000.
- Donna Tjandra, Yifei He, and Jenna Wiens. A hierarchical approach to multi-event survival analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 591–599, 2021.
- Stéphanie Van Der Pas, Rob Nelissen, and Marta Fiocco. Different competing risks models for different questions may give similar results in arthroplasty registers in the presence of few events: illustrated with 138,234 hip (124,560 patients) and 139,070 knee (125,213 patients) replacements from the dutch arthroplasty register. *Acta Orthopaedica*, 89(2):145–151, 2018.
- Bob JH van Kempen, Bart S Ferket, Maryam Kavousi, Maarten JG Leening, Ewout W Steyerberg, M Arfan Ikram, Jacqueline CM Wittteman, Albert Hofman, Oscar H Franco, and MG Myriam Hunink. Performance of framingham cardiovascular disease (cvd) predictions in the rotterdam study taking into account competing risks and disentangling cvd into coronary heart disease (chd) and stroke. *International journal of cardiology*, 171(3):413–418, 2014.
- Hilary K Wall, Matthew D Ritchey, Cathleen Gillespie, John D Omura, Ahmed Jamal, and

Mary G George. Vital signs: prevalence of key cardiovascular disease risk factors for million hearts 2022—united states, 2011–2016. *Morbidity and Mortality Weekly Report*, 67(35):983, 2018.

Ping Wang, Yan Li, and Chandan K Reddy. Machine learning for survival analysis: A survey. *ACM Computing Surveys (CSUR)*, 51(6):1–36, 2019.

Peter WF Wilson, Ralph B D’Agostino, Daniel Levy, Albert M Belanger, Halit Silbershatz, and William B Kannel. Prediction of coronary heart disease using risk factor categories. *Circulation*, 97(18):1837–1847, 1998.