

LEARNING GENOMIC REPRESENTATIONS TO PREDICT CLINICAL OUTCOMES IN CANCER

Safoora Yousefi[†] Congzheng Song[†] Nelson Nauata[‡] Lee Cooper^{‡,*}

[†]Department of Mathematics & Computer Science Emory University

[‡]Department of Biomedical Informatics Emory University School of Medicine

*Department of Biomedical Engineering Georgia Institute of Technology

{safoora.yousefi, congzheng.song, nnauata, lee.cooper}@emory.edu

ABSTRACT

Genomics are rapidly transforming medical practice and basic biomedical research, providing insights into disease mechanisms and improving therapeutic strategies, particularly in cancer. The ability to predict the future course of a patient's disease from high-dimensional genomic profiling will be essential in realizing the promise of genomic medicine, but presents significant challenges for state-of-the-art survival analysis methods. In this abstract we present an investigation in learning genomic representations with neural networks to predict patient survival in cancer. We demonstrate the advantages of this approach over existing survival analysis methods using brain tumor data.

1 INTRODUCTION

Genomics provide a window into the complex molecular workings of disease. In the treatment of cancer, genomic analysis of a tissue biopsy can reveal specific molecular vulnerabilities that can be matched to targeted therapies, or to prognosticate the future behavior of a patient's disease and expected survival in order to better inform clinical interventions including surgery and radiation therapy. Although genomic analysis generates rich high-dimensional signals that contain hundreds to hundreds-of-thousands of variables, typically only several variables are used for prognostication for any given cancer type. Typically, these variables are used to assign patients into discrete disease classes or "subtypes" that associate with response to specific therapies, or with varying degrees of disease aggressiveness. Learning the underlying latent prognostic variables from high-dimensional genomic profiles can extract additional prognostic value from unused variables, and is critical in realizing the promise of genomic medicine. This problem presents significant challenges, ranging from the familiar "large p small N ", to how to adapt developments in the machine learning domain to the analysis of time-to-event survival data.

In this abstract we present an investigation in building survival prediction neural networks to learn representations from genomic data for survival prediction. We use backpropagation to train neural networks to maximize the Cox proportional hazards likelihood of time-to-event data, and apply these predictive models to molecular profiles of brain tumor patients from The Cancer Genome Atlas where survival ranges from 6 months to 10+ years. We compare our methods to state-of-the-art survival analysis algorithms based on elastic-net regularization of Cox hazard models and random forest based methods, and demonstrate improvements in survival prediction accuracy for neural network approaches.

2 BACKGROUND AND RELATED WORK

2.1 HAZARD MODELS AND LIKELIHOOD FUNCTIONS

Survival analysis involves predicting the time to some event of interest, which in cancer is often death or progression of disease. It differs from ordinary regression due to *incomplete followup*, where a death or relapse event is not observed at or before the final encounter with the patient. These censored observations provide critical information, and often represent an important population of long-term survivors or treatment responders that are very important to incorporate into the model. The most commonly used regression approach to survival analysis is the Cox proportional hazards

model proposed by Cox (1972). At time t , the hazard for a sample with covariates x is given by the following hazard function:

$$\lambda(t|x) = \lambda_0(t)e^{\beta x}, \tag{1}$$

where $\lambda_0(t)$ is baseline hazard. The hazard function is an exponential linear function of the covariates x and model coefficients β , with the effect of any covariate $x^{(i)}$ assumed to be the same over time. Since the evaluation criteria of the models in this paper is based on ranking predicted survival times, the baseline hazard is left unspecified and we maximize the partial likelihood function during training:

$$l(\beta, X) = - \sum_U \left(X_i \beta - \log \sum_{j \in R_i} e^{X_j \beta} \right) \tag{2}$$

where U is the set of all uncensored patients, and R_i is the set of patients whose time of death or last follow-up is later than time of death of i .

We measured model performance using *concordance index* (CI) that captures the rank correlation of predicted and actual survival. Denoting the i th patient with X_i and the set of all patients with X , where t_i represents either the time of death or the time of last follow-up of the i th patient, CI was calculated in the following way:

$$CI(\beta, X) = \sum_P \frac{I(i,j)}{|P|} \tag{3}$$

$$I(i, j) = \begin{cases} 1, & \text{if } Risk_j > Risk_i \text{ and } t_j > t_i \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

Where P is the set of orderable pairs. A pair of samples (X_i, X_j) is orderable if either the event is observed for both X_i and X_j , or X_j is censored and $t_j > t_i$.

Intuitively, CI measures the pairwise agreement of the prognostic scores $Risk_i, Risk_j$ predicted by the model and the actual time of death for all orderable pairs. Attempts have been made to propose differentiable versions of CI based on sigmoid and exponential functions and optimize it directly. But recent studies show that optimizing the cox partial likelihood is equivalent to optimizing CI (Steck et al., 2008).

2.2 RELATED WORKS

Regularization techniques have been proposed for feature selection in survival analysis of high dimensional data (Zou & Hastie, 2005). Efforts have been made to introduce successful machine learning algorithms such as random forests to survival analysis (Ishwaran et al., 2008). Deep learning techniques have been employed for cancer diagnosis using genomic data and medical images, such as Fakoor et al. (2013) and Esteva et al. To the best of our knowledge, representation learning techniques have not been applied to survival prediction from genomic data, and the previous work investigating neural networks for survival analysis dealt with low dimensional data and different cost functions (Lisboa et al., 2003).

3 SURVIVAL PREDICTION NEURAL NETWORK

3.1 PRETRAINING AND FINE-TUNING

In this work we trained an autoencoder to represent genomic data and fine tune this representation using partial log Cox likelihood. In training, we employ stacked denoising autoencoders proposed in Vincent et al. (2008). We train the auto-encoders using 183-dimensional genomic features, then we add a risk prediction output layer as shown in Figure ??.

We use survival times and censoring status to calculate the Cox partial log likelihood given by equation 2 and differentiate it with respect to X :

$$\frac{\partial l(\beta, X)}{\partial X_i} = c_i \beta - \sum_{j \in U: i \in R_j} \frac{e^{X_i \beta}}{\sum_{k \in R_j} e^{X_k \beta}} \tag{5}$$

where c_i is 1 if sample i is not censored, and is 0 otherwise, and β denotes the parameters of the risk prediction layer. This derivative is then back-propagated through the network to fine tune the learned representation specifically for the task of survival analysis.

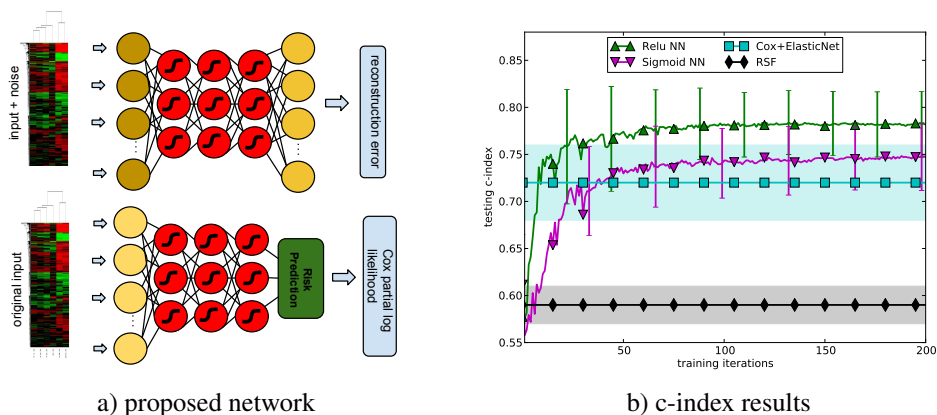


Figure 1: a) Survival network model. b) Comparison of the proposed survival prediction neural network to two competing methods: regularized Cox regression and Random Survival Forest(RSF). Average testing CI trends are shown for neural networks with two different activation functions: sigmoid and rectified linear units. The error bars and shaded areas indicate standard deviation of CI over 10 cross validation sets

3.2 MODEL SELECTION

The training of a neural network involves many hyperparameters: type of nonlinearities used, number of layers, number of hidden units in each layer, learning rates for pretraining and fine tuning and regularization parameters. Since this is the first work where deep neural networks are used to address survival analysis, we could not look at existing literature for a conventional choice of hyperparameters. Unlike areas such as image classification, no rule of thumb has been developed for setting hyper-parameters in survival analysis. Therefore we employed bayesian optimization (Martinez-Cantin, 2014) with Gaussian prior to decrease the number of objective function evaluations needed to reach a decent choice of hyperparameters. More shallow networks demonstrate superior performance over deeper architectures in our experiments. This could be justified considering the small number of training samples (628) and the scarcity of labels within the available samples. Our average choice of configuration is 2 fully connected layers of 250 hidden units each. On average, we use a learning rate .001 for pre-training and .0009 for fine-tuning.

3.3 EVALUATION

Due to the small size of available training data, performance of the model might considerably depend on the partitioning of the data into testing and training. To mitigate this, we sampled %70 of the data set 10 times without replacement to have 10 different training sets. In every training episode, we used half of the remaining %30 of data for model selection and the other half for model assessment. The reported CI is averaged over the 10 experiments to represent the generalization error of the model. The exact same setting was used for hyper-parameter tuning and assessment for competing methods. We picked the learning rate and elastic-net mixture coefficient for regularized Cox regression (Hastie & Qian (2014)) based on performance on the same validation sets we used for the neural networks. We tuned number of trees, leaf size, and number of split points for random survival forest in the same fashion. Our experiments reveal that Random Survival Forests do not adapt well to high dimensionality and are markedly outperformed by survival neural networks (See Figure 1). Neural networks also achieve %5 absolute improvement over regularized Cox regression with ReLU activation and %3 with sigmoid activation.

ACKNOWLEDGMENTS

This work was supported by US Public Health Service National Institutes of Health (NIH) grants K22LM011576-03 and U24CA194362-01.

REFERENCES

- David R Cox. Regression models and life tables (with discussion). *Journal of the Royal Statistical Society*, 34:187–220, 1972.
- Andre Esteva, Brett Kuprel, and Sebastian Thrun. Deep networks for early stage skin disease and skin cancer classification.
- Rasool Fakoor, Faisal Ladhak, Azade Nazi, and Manfred Huber. Using deep learning to enhance cancer diagnosis and classification. In *The 30th International Conference on Machine Learning (ICML 2013), WHEALTH workshop*, 2013.
- Trevor Hastie and Junyang Qian. Glmnet vignette. Technical report, Technical report, Stanford, 2014.
- Hemant Ishwaran, Udaya B Kogalur, Eugene H Blackstone, and Michael S Lauer. Random survival forests. *The annals of applied statistics*, pp. 841–860, 2008.
- Paulo JG Lisboa, H Wong, P Harris, and Ric Swindell. A bayesian neural network approach for modelling censored data with an application to prognosis after surgery for breast cancer. *Artificial intelligence in medicine*, 28(1):1–25, 2003.
- Ruben Martinez-Cantin. Bayesopt: A bayesian optimization library for nonlinear optimization, experimental design and bandits. *The Journal of Machine Learning Research*, 15(1):3735–3739, 2014.
- Harald Steck, Balaji Krishnapuram, Cary Dehing-oberije, Philippe Lambin, and Vikas C Raykar. On ranking in survival analysis: Bounds on the concordance index. In *Advances in neural information processing systems*, pp. 1209–1216, 2008.
- Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pp. 1096–1103. ACM, 2008.
- Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.