

Health Outcomes in Mid-Ages: Multistate time to event Statistical Models versus Long Short Term Memory (LSTM) Recurrent Neural Network (RNN) Models*

Lakshmi K. Raut¹

¹Visiting Fellow, University of Chicago

¹lakshmiraut@gmail.com

1/19/23

Abstract

Health outcomes such as serious diseases, disability and death develop sequentially over one's life span. Genetic and epigenetic factors, and health related behaviors throughout one's life condition those health outcomes. Some individuals develop one or more diseases and their health deteriorate faster, worsening their quality of lives and survival probabilities. The models in health economics literature formulate progression of health outcomes over the life span using the multistate time to event framework, and estimate the effects of the above mitigating factors on probability of transitions from one health state to another. The sequential nature of health progression is captured in the Markovian structure. Markov chain models have short memory, i.e., these models assume that given the current health outcome, the past does not influence the probability of transition to another health outcome. Many chronic diseases such as cancers and heart diseases manifest as a result of long lagged past health behaviors. Markov chain models are limited in capturing these effects. More recently, long short term memory (LSTM) recurrent neural network (RNN) models are developed in the machine (deep) learning literature that keep track of important features from the past inputs in its memory cells, which the model determine important in the 'context' of the future outcomes. These models are generally applied to natural language processing, creating audio streams, and video subtitling. In this paper, I adapt the existing LSTM-RNN models for the prediction of sequential health outcomes in mid-ages. I first compare the merits and shortcomings of these two approaches and then use the Health and Retirement Study (HRS) data to compare their performances in predicting sequential health outcomes.

JEL Classifications: I12, C41, C51.

Keywords: Forecasting health outcomes, multistate time-to-event model, long short term

*Draft prepared for presentation at the ASSA 2020 Annual Meeting, January 3 - 5, 2020, San Diego, California, USA. I got many insightful comments from Han Altae-Tran at MIT, John Phillips at NIH, and comments from the discussant and the audience. Thanks. **Disclaimer:** The views, thoughts, and opinions expressed in the paper belong solely to the author, and do not necessarily represent the views of any institution, other group or individual.

memory (LSTM) recurrent neural network (RNN) model.

One-liner: The paper compares performances of a statistical multistate time to event model and a long short term memory (LSTM) recurrent neural network (RNN) model in forecasting sequentially the health outcomes over the life span.

Contents

List of Figures	2
List of Tables	3
1 Introduction	3
2 Statistical multi-state model of health outcomes process	8
2.1 Aalen-Johansen Estimator of Transition Probabilities	13
3 The data set and the variables	16
3.1 Variables	17
3.2 Statistical Results	18
4 Long-short memory (LSTM) recurrent neural network (RNN) model of health outcomes process	22
4.1 The basics of neural network	22
4.2 Previous neural network models of survival analysis	25
4.3 Formulation within Recurrent Neural Network with LSTM	26
4.4 LSTM memory cell	28
5 Prediction and performance of the model	29
6 Conclusion	31
Appendix	32
MLP learning algorithm	32
Bibliography	33

List of Figures

1	Path diagram of health trajectors.	10
2	Transition probabilities (a) from normal health state, (b) from one-or-more diseased health state	19
3	MLP architechture.	24
4	Showing a simple recurrent neural, with layer 0 is the input layer, layer 1 is the recurrent hidden layer, layer 2 is another pure hidden layer, layer 3 is the output layer, for each time point (t), layers are denoted as superscript, and the subscript in parenthesis is time.	26
5	Showing a typical memory cell at time-step t	28
6	Means and mean \pm standard deviation intervals of the predicted cumulative incidence rates of disability and death in the test sample	30

List of Tables

1	Summary of the health status of the individuals in the sample over the survey years.	17
2	Estimated transition probabilities for transitions ioj by duration of stay in state 1 and 2.	18
3	Estimates of Cox regression models separately for each transition with demographic variables only	20
4	Estimates of Cox regression models separately for each transition with health measures	21
5	Average of the Predicted cumulative incidence rates of disability and death for three methods for the test sample.	30

1 Introduction

The path through various health outcomes such as serious diseases, disability and death develop sequentially over one's life span. Many factors such as genetic and epigenetic factors, and health related behaviors throughout one's life condition those health outcomes. Some individuals develop one or more diseases and their health deteriorate faster, worsening their quality of lives and survival probabilities, while others have slower aging process. The biology of aging process determines how those factors affect aging at the cellular level. In this paper, with insights from the biology literature, I use statistical and neural network models for predicting and estimating the transition

probabilities between health states and the times they spend in a health state before transiting to another health state for individuals in their mid ages. While the framework I develop in the paper is applicable to many other health outcomes such as incidence of cardiovascular and immune diseases, I focus on one health outcome of interest—any of the disabilities that qualifies for a public disability insurance program—with death as a competing risk health outcome. Death is a competing risk health outcome in the sense that once dead, one cannot be at risk of disability, or as a matter of fact, any other health outcome of interest.

The main issues I address in this paper are: What kind of aging process and what conditioning factors for it the biology literature recommends? Which type of model—a multistate statistical model, a feedforward multilayer perceptron (MLP) type of neural network model, or a long short memory (LSTM) recurrent neural network (RNN) model—that is better suited for prediction and estimation of transition probabilities among health states, taking into account the effect of various covariates (the conditioning factors)? Which model can estimate the effect of various conditioning factors and their relative importance on the type of health trajectory one follows?

I will not get into the details of the biomedical literature on these issues. Here I will point out the main insights from my reviews elsewhere, ([Raut, 2019a](#); [Raut, 2019b](#)). Similar to the literature of behavioral genetics of personality and intelligence, the *nature-nurture* controversy exists in the health literature: Is it all nature (i.e., all genetics or genome) or is it all nurture (i.e., all epigenetics or epigenome modulated by the environment and health related individual behaviors) that determines the progression of health over the life span of an individual? The consensus so far is that it is neither the nature nor the nurture, it is a combination of the two that determines health developments over one's life-span. The research so far found that certain genetic make-ups (i.e., certain sequences of DNA) predispose one to certain diseases, (([Barondes, 1999](#); [Khoury et al., 2009](#); [Bookman et al., 2011](#))), but the epigenetic inputs—especially at the very early stage of life, i.e. in the womb, but not the least at later stages of life—are also very important determinants of life expectancy and quality of life. The biomedical research so far has not found genes that are responsible for aging and age related diseases, leading to early disability and mortality. The twenty-first century biomedical research emphasizes more on the epigenetic factors than the genetic factors to explain the pattern of health developments over the life-span.

At the cellular level, aging means cellular senescence—i.e., after a certain number of cell divisions, it stops dividing or have defective replications, causing tissues or organs to increasingly deteriorate over time. Senescence leads to incidence of degenerative diseases. It is generally observed that women live longer than men and those with better life styles in terms of smoking, exercising

and diets delay the aging process (for evidence, see (Blair et al., 1989; Vaupel, 2010; Austad and Fischer, 2016; Zarulli et al., 2018)). This line of biological inquiry led to explore the (cellular) molecular mechanism of aging process and to find biomarkers of aging that can be used to diagnose, monitor, and improve the age related physiological decline and disease. A good indicator of the aging process at the cellular level is the rate of decay in the telomere length. Telomeres are the caps at the end of chromosomes in a DNA sequence. They look like the plastic caps at the end of shoelaces. The main function of telomeres is to protect cells preserving the genetic content within each chromosome during cell divisions. Unfortunately, the telomere length shortens in the course of each cycle of chromosomal replication during cell division, reaching the Hayflick limit (about 40 to 60 cell divisions, Hayflick (1965)) with a critically short telomere length, after which the cells stop dividing or divide with chromosomal abnormalities. The rate of shortening of the telomere length is modulated by telomerase enzyme. Why the rate of decay in telomere length varies for individuals is an active area of biomedical research and the mechanism for it is not yet fully understood. Many studies find that higher stress of any kind— psychological, financial, social and chemical—is strongly associated with higher oxidative stress, lower level of telomerase enzyme, and shorter telomere length. Furthermore, shorter telomere length is associated with health related phenotypes of poorer health and higher risks for cardiovascular and immune diseases (see, (Epel et al., 2004; Shalev, Entringer, et al., 2013; DiLoreto and Murphy, 2015; Shalev and Belsky, 2016; Simons et al., 2016)).

More recently emerged second line of biomedical research on aging and aging related diseases explores the epigenetic (which literally means on top of genetic) mechanism for these life-cycle processes. (See for instance, (Alisch et al., 2012; Barres and Zierath, 2011; Boks et al., 2009; Esteller, 2008; Hannum et al., 2013; Horvath, 2013)).

The above literature emphasizes that aging and age related diseases are associated with shortening of telomere length and changes in global methylation, and that stress, smoking, drinking, chemical misuse, physical exercising, and diet are important modulators for these changes. The question remains, what are the critical periods or the developmental milestones in life-cycle that program the motions of health developments over the life-span of an individual?

Research along this line began with the striking findings of (Barker, 1990; Barker, 1998) and later of Gluckman et al. (2008). They found strong associations between birth weight and many later life chronic diseases, including hypertension, coronary artery disease, type 2 diabetes, and osteoporosis. Many other studies find that much of health developments in later life is determined very early in life, during the prenatal period right after conception, i.e. in the womb. Sometimes it is

said in social sciences that inequality begins in the womb. The effect of an environmental stress in the womb on later life diseases and developmental outcomes is known as *programming*. [Gluckman et al. \(2008\)](#) observes that “like the long latency period between an environmental trigger and the onset of certain cancers, the etiology of many later life diseases such as cardiovascular disease, metabolic disease, or osteoporosis originate as early as in the intrauterine development and the influence of environments that created by the mother.’’ For more empirical evidence for the developmental origin of later life diseases, see ([Barker, 2007](#); [Thornburg et al., 2010](#)). The papers by ([Kanherkar et al., 2014](#); [Barbara et al., 2017](#)) provide detailed descriptions of the biological process of development of life and health, starting from conception. They explain how the global DNA demethylation of the fertilized egg right after conception creates an epigenetic “clean slate” to start a new life, followed by rapid remethylation to reprogram the maternal and paternal genomes to create epigenetic configurations in the fetus which rapidly produce specialized cells of the body with cell divisions. The environment provided in mother’s womb during those times has long-term effects on the child’s later cognitive and other health developments. While inputs at early milestone ages are important for later age health, healthy living and good healthcare are still important for maintaining health in mid ages.

Studies in social sciences find that low socio-economic status (SES) are associated with inflammation, metabolic dysregulation, and various chronic and age-related diseases such as type 2 diabetes, coronary heart disease, stroke, and dementia, and that low SES create epigenetic changes in individuals that lead to faster biological aging even after controlling for health-related behaviors such as diet, exercise, smoking, alcohol consumption, or having health insurance, see for evidence, [Simons et al. \(2016\)](#). The study by [Karakus and Patton \(2011\)](#) uses the Health and Retirement Study data and after controlling for education, race, income, health risk indicators like BMI and smoking, functional limitations like gross motor index, health limitations for work, and income, they find depression at baseline leads to significantly higher risk for developing diabetes, heart problems, and arthritis and no significant effect on developing cancer during the 12 years follow-up period. [Renna \(2008\)](#) uses National Longitudinal Survey of Youth data to find no significant effect of alcohol use on labor market outcomes such as on earnings or hours of work. [Seib et al. \(2014\)](#) collected data on a sample of older women in Australia and found that severe traumatic life events create strong stress levels that influence them to have unhealthy living and diet measured by BMI and develop stronger and earlier health problems. [Conti et al. \(2009\)](#) utilize the CES-D data in the Health and Retirement Study dataset to construct a measure of depression, and find that depression of men and women have significant negative effect on employment status, early retirement, and application for DI/SSI benefits.

Using insights from the above literature, I formulate a finite state continuous time stochastic process model of disablement process. I postulate that as individuals age, the homeostatic regulatory mechanism that controls physiological systems—respiratory, cardiovascular, neuroendocrine, immune, and metabolic—becomes more and more fragile in its ability to face internal and external stressors, leading to early occurrence of disease, disability and death. I use available bio-markers (such as BMI, CES-D, cognition) and health related behaviors such as smoking, and vigorous exercising along the life-course to explain how they affect the risk of chronic diseases, disabilities and death.

A multistate stochastic process framework is useful to study the effects of various covariates—the covariates that may be different for different intermediate health states—on the risks of disability and death.

The statistical models for estimation of transition probabilities are based on Markov processes and assume that transition intensities of these processes have Cox proportional hazard specifications, see for instance (([Aalen and Johansen, 1978](#); [Andersen, O. Borgan, et al., 1993](#); [Andersen and Perme, 2008](#); [Crowther and Lambert, 2017](#); [Fleming, 1978](#))). Markov models have short memories and proportionality assumption imposes serious limitations on the structure of the model, which could be far from the functional form of the true data generating process. Neural network models relax these limitations.

It is known that a feedforward multilayer perceptron (MLP), also known as a feedforward deep learning model, with a sufficient number of neurons in the hidden layer can approximate any function as closely as desired. That is, an MLP is one of the best ‘universal function approximator’ ([Hornik et al. \(1989\)](#)). A few papers—([Faraggi and Simon, 1995](#); [Biganzoli et al., 1998](#); [Katzman et al., 2018](#); [Lee et al., 2018](#); [Ranganath et al., 2016](#))—used feedforward MLP networks to compute the survival probabilities when there is only one possible transition between two health states—alive and death—with the exception of [Lee et al. \(2018\)](#) who studied competing risks, by breaking death into various causes of death. [Katzman et al. \(2018\)](#) introduced more general non-linearity of the covariate effects, but still kept the Cox proportionality assumption. [Ranganath et al. \(2016\)](#) assumed parametric form for the baseline hazard function as compared to the non-parametric form in Cox model, but they made the covariate effects nonlinear. [Ren et al. \(2019\)](#) consider a recurrent neural network, but the covariates are time-fixed at the initial time step. They also restricted to only one transition, i.e., a two-state model. None of these models deals with sequential framework where new information arise with time steps and update the previously estimated transition probability estimates. In these models, all the inputs from the past, present and the future times in the sample determine current probabilities. These models, with the exception of [Ren et al. \(2019\)](#)

although with other serious limitations, have no ways to store information learned from the past inputs. After training these models, when new data arrive, these models cannot use this new data to update the predicted probabilities without losing information in the early periods.

A recurrent neural network (RNN) uses feedback connections or self connection of neurons in the hidden layer, and thus is capable of storing important information learned in the past in these recurrent neurons. Like an MLP is a universal function approximator, an RNN has the similar nice property that with sufficiently large number of hidden recurrent neurons, an RNN can approximate any sequence-to-sequence mapping ((Graves et al., 2014; Hammer, 2000; Siegelmann and Sontag, 1992)). These models have shared weights between time-steps and in the input and output layers, as a result when new data arrive after training the model, it can use all the past important information learned from the past to this new data point and predict the future probabilities in the light of this new data. Since training such models involve computation of gradients using backpropagation through time, it involves multiplication of numbers less than one many times, leading to vanishing gradient problem. In these scenarios, it cannot keep useful information in memory from the long time back. Overcoming these problems led to a few modifications of the RNN framework. The most successful of them is the long short memory (LSTM) RNN model introduced by Hochreiter and Schmidhuber (1997). For more on LSTM-RNN models, see Graves (2012). I use this LSTM-RNN model for prediction of time-to-event probabilities of health outcomes. Another problem is with the training data size. To obtain good predictive performance, these models require a large number of training examples. In drug discovery problems or with surveys or lab experiments, obtaining large number of examples is costly. To overcome small data problem, (Altae-Tran, 2016; Altae-Tran et al., 2017) introduced further refinement of the LSTM-RNN framework. I do not adopt such modifications. In this paper, I use the original LSTM-RNN model specified in Graves (2012) and implemented in Keras module of Tensorflow 2.0 for Python package.

The rest of the paper is organized as follows. In Section 2, I describe the econometric specification of the multistate stochastic process and describe estimation issues. In Section 3, I describe the Health and Retirement Study data set and the variables that I use in estimation in both frameworks. In Section 4, I describe the recurrent neural network with LSTM memory cells that I use to extend the multistate Cox model of health outcomes process. In Section 5, I describe the performance criteria that I use to compare the performance of the models used in this study. Section 6 concludes the paper.

2 Statistical multi-state model of health outcomes process

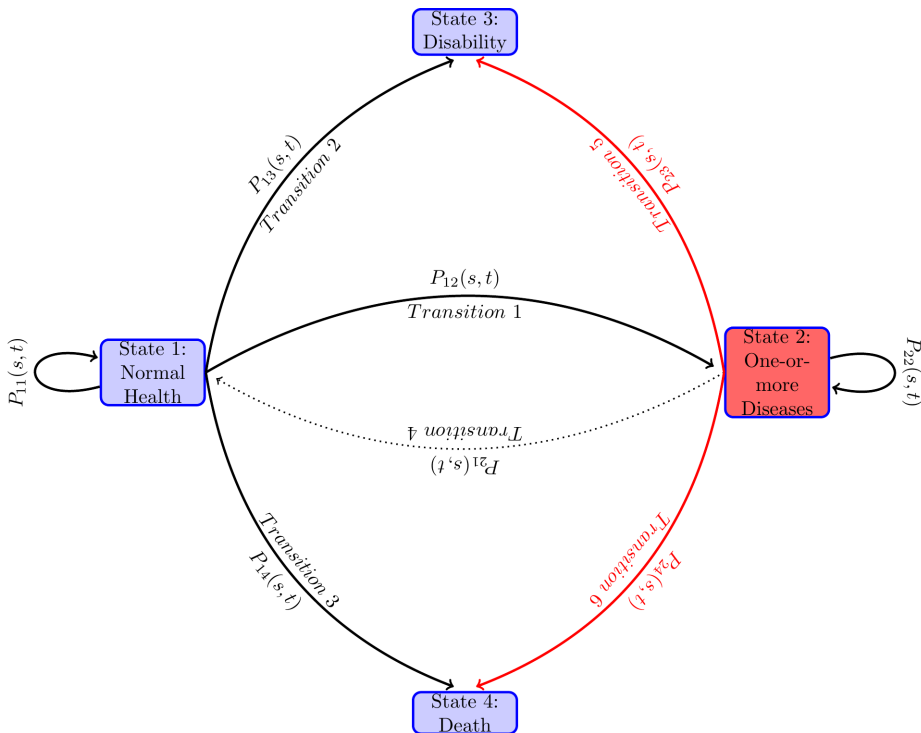
The goal is to formulate and then estimate an econometric model of paths to enter disability rolls. An individual can be on the disability rolls if the individual has a qualifying disability before reaching age 65 and has not died before applying for disability benefits. I assume that an individual's getting on the disability rolls is a terminal event, i.e., the individual does not move to normal or diseased health states. After reaching this state, the individual is not followed any further. A competing risk for getting on the rolls is death before age 65. This is a competing risk because an individual cannot be at risk for disability enrollment if the individual is already dead and thus not at risk to get on the disability rolls. The individual is not followed after the event of death because our primary interest is the event of getting on to the disability roll. In the technical terms defined below, we treat the health states—disability and death—as absorbing states, i.e., once in that health state, an individual remains in that health state and removed from the sample for later considerations. An individual can be in normal health and then become disable or die before becoming disabled or may first become diseased with one or more diseases and again from that health state become disabled or die before becoming disabled. Various factors affect individual risks of various transitions of health states and the time they stay in each health state along the life-span. Both, in turn, determine the timing of getting on to the disability rolls.

I model the paths through various health states that individuals follow along their life-spans as a continuous-time finite-state Markov process $X(t), t \in T$, where at each time point t during the study period T , the random variable $X(t)$ takes a value from a finite number of health states in S . In the present study, we take $T = [0, 7]$, treating age 51 as time period $t = 0$ and age 65 as time period $t = 7$. The unit of time is 2 years, as HRS collected data every two years. The state space S contains states, 1 = “healthy or normal health”, 2 = “ill with one or more chronic disease”, 3 = “disabled with DI-or SSI-qualifying disability” and 4 = “Death”. Sometimes I will use $S = \{h, i, d, D\}$ in place of $\{1, 2, 3, 4\}$.

Typically in the study period, an individual along the path to disability or death before age 65 may be in the normal health state for a length of time, and then moves to another health state, say diseased health state, and remain there for some time, and then jump to the health state of disability or to death, or reach 65 and censored. There are many possible paths that an individual can follow. Even when the health states they pass through are the same, the duration of stay in each health state (also known as the *waiting time* in stochastic process literature) could vary. Each configuration of visited states and the waiting times in those states constitute one path. When time is continuous, an uncountably large number of paths are possible. From the diagram below one can see various

paths that an individual may follow during the study period. The focus of the paper is to study the probabilities of various transitions and the duration of stay in each health state.

Figure 1: Path diagram of health trajectories.



Let the transition probabilities of our Markov process $X(t)$ be given by

$$P_{hj}(s,t) = \text{Prob}(X(t) = j | X(s) = h), \quad (1)$$

for all $h, j \in S$, and $s, t \in T$, $s \leq t$. Denote the matrix of transition probabilities by

$$P(s,t) \equiv (P_{hj}(s,t))_{h,j=1\dots p}.$$

An individual may be in any of the health states in S at time t , the probability of which, known as the *occupation probability*, depends on the occupation of the previous health states. and the transition probabilities among health states. Let $\pi_j(t)$ be the occupation probability of an individual in health state j at time t . Occupation probability can be also viewed as the proportion of population of age t who are in health state j . Denote all the occupation probabilities as a column vector $\pi(t) \equiv (\pi_j(t)), j \in S$. Then the occupation probabilities move over time recursively as follows,

$$\pi(t) = \pi'(s)P(s,t), 0 \leq s < t.$$

The goal is to estimate for each time $s, s \in T$ the transition probability matrices $P(s, t), t > s, t \in T$, taking into account the effects of (i.e., conditioning on) the past materialized values of the determining factors, that is covariates, and the stochastic process $X(\cdot)$ up to time s .

It is known that the transition probabilities of a Markov process satisfies the Chapman-Kolmogorov equation

$$P(s, t) = P(s, u) \cdot P(u, t), \text{ for all } s, u, t \in T \text{ with } s < u < t \quad (2)$$

I will use the above to point out how the transition probabilities are parameterized and statistically estimated by the Aalen-Johnson non-parametric estimation method that plugs in the Nelson-Aalen estimates of the integrates hazards for each transition separately in an unified way for both continuous time and discrete time Markov processes.

I assume that the transition probabilities $P(s, t), s, t \in T, s < t$ are absolutely continuous in s and t . A *transition intensity*— also known as the *hazard rate* in the survival analysis literature when there is only one possible transition (alive to death), and as the *cause-specific hazard rate* in the competing risk analysis when death is split into various causes of death¹ —of the health process X_t from health state h to health state j at time t is the derivative

$$\begin{aligned} \lambda_{hj}(t) &= \lim_{\Delta t \rightarrow 0} \frac{P_{hj}(t, t + \Delta t) - P_{hj}(t, t)}{\Delta t}, \text{ for } j \in S, \text{ which for } j \neq h \text{ becomes} \\ &= \lim_{\Delta t \rightarrow 0} \frac{P_{hj}(t, t + \Delta t)}{\Delta t}, \text{ and for } j = h \text{ becomes} \\ \lambda_{hh}(t) &= \lim_{\Delta t \rightarrow 0} \frac{P_{hh}(t, t + \Delta t) - 1}{\Delta t} \\ &= - \lim_{\Delta t \rightarrow 0} \frac{\sum_{j \neq h} P_{hj}(t, t + \Delta t)}{\Delta t} \\ &= - \sum_{j \neq h} \lambda_{hj}(t) \end{aligned}$$

For absorbing states $h = 3, 4$, the transition intensities $\lambda_{hj}(t) = 0$, for all $j, j \in S$. Denote the matrix of transition intensities by

¹See for instance, [Raut \(2017\)](#) for a competing risk analysis in a similar context using the SSA Administrative data and compare that with the present framework.

$$\Gamma(t) = \begin{pmatrix} -(\lambda_{12}(t) + \lambda_{13}(t) + \lambda_{14}(t)) & \lambda_{12}(t) & \lambda_{13}(t) & \lambda_{14}(t) \\ 0 & -(\lambda_{23}(t) + \lambda_{24}(t)) & \lambda_{23}(t) & \lambda_{24}(t) \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (3)$$

Every Markov stochastic process has an associated intensity process. In the theory of stochastic processes, it is known that given an intensity function $\Gamma(t)$, there exists a continuous time Markov chain satisfying the Chapman-Kolmogorov Eq. (2) and conversely, given a Markov chain with transition probabilities $P(s, t)$, $s, t \in T$, $s < t$, that satisfies the Chapman-Kolmogorov Eq. (2), one can derive the transition intensities of the form in Eq. (3). For this reason, the intensity function $\Gamma(t)$ is also referred as *infinitesimal generator of the Markov process*. While $P_{hj}(s, t)$ is an unconditional probability, the transition intensity or the hazard rate $\lambda_{hj}(t)\Delta t$ is the conditional (instantaneous) probability. More specifically, $\lambda_{hj}(t)\Delta t$ is the probability that an individual experiencing event j in a very small interval of time $[t, t + \Delta t)$ given that he has been in state h at time t . This conditional probability may depend on time t and other characteristics up to time t .

In statistical and neural network models, the dependence of transition probabilities on individual characteristics is generally done through parametric or semi-parametric specification of the transition intensity functions $\Gamma(t)$. One then estimates the transition probabilities from the non-parametric or semi-parametric estimates of the integrated transition functions. An *integrated transition intensity function* $\Lambda_{hj}(t)$ for a transition $h \rightarrow j$ is defined by $\Lambda_{hj}(t) = \int_0^t \lambda_{hj}(u) du$. Just like for a continuous random variable, it is easier to estimate its cumulative density function non-parametrically than its density function, for the time-to-event data with censoring, it is easier to estimate the integrated hazard function than the intensity function. While for the discrete case this problem does not arise, estimation of transition probabilities via nonparametric estimates of the integrated hazard function is an unified approach encompassing both discrete time and continuous time data. I follow this strategy.

To explain and gain insights into this estimation strategy, denote the matrix of all the cumulative hazard functions as $\Lambda(t) = (\Lambda_{hj}(t))_{h,j=1,2,3,4}$, and the matrix of the derivatives of the component functions by $d\Lambda(t)$. Let the time interval $[s, t]$ is subdivided into a partition of m sub-intervals with cut-off points $s = t_0 < t_1 < \dots < t_m = t$. Denote the partition by $P(m)$. Denote the largest size of the sub-intervals by $|P(m)| \equiv \max\{|t_i - t_{i-1}|, i = 1, \dots, m\}$. Applying repeatedly the Chapman-Kolmogorov Eq. (??) on the sub-intervals of the partition, we have

$$P(s, t) = P(t_0, t_1) \cdot P(t_1, t_2) \cdot \dots \cdot P(t_{m-1}, t_m) = \prod_{i=1}^m P(t_{i-1}, t_i) \quad (4)$$

Note that as $|t_{i-1} - t_i| \rightarrow 0$, the transition probability matrix $P(t_{i-1}, t_i) \rightarrow P(t_{i-1}, (t_{i-1} + dt)) = I + \Gamma(t)dt$.² As $|P(m)| \rightarrow 0$, the right hand side of Eq. (??) converges to a matrix, known as the *product integral*³ of the integrated hazard functions $\Lambda(s, t)$, denoted as $\prod_s^t (I + d\Lambda(u))$. Or in other words, the transition probabilities of a stochastic process parameterized via an intensity process is given by the product integral of integrated hazard function.

$$P(s, t) = \prod_s^t (I + d\Lambda(u)). \quad (5)$$

The above product-integral solution is a generalization of the Kaplan-Meier (Kaplan and Meier (1958)) product-limit formula for the survival function in survival analysis. The product integral formula unifies both discrete time and continuous time Markov processes, and is an extremely useful apparatus for statistical analysis of Markov processes.

The most widely used statistical procedure to estimate the transition probabilities $P(s, t)$, $s, t \in T, s < t$ is to plug in an estimate $\Lambda(u)$ in Eq. (5). The effect of covariates is incorporated by conditioning the transition intensity functions $\Gamma(t; X(t))$ on the covariates process $X(t)$. There are many ways to get these estimates. I will follow two approaches in this paper: First, I will explore the more widely used non-parametric Aalen-Johnson-Fleming method via Nelson Aalen estimates for each-component of the $\hat{\Lambda}(u; X(u))$ with Cox proportional hazard model to incorporate the time-varying covariate effects in the next sub-section. Second, the Neural network approach explored in a later section.

2.1 Aalen-Johansen Estimator of Transition Probabilities

Most widely used statistical procedure incorporates the time-varying covariates for the transition probabilities by specifying a semi-parametric functional forms for the intensity hazard functions

$$\lambda_{hj}(t; X(t)) = \lambda_{hj}^0(t) e^{\beta_{hj}' X(t)} \quad (6)$$

²From definition of transition intensities above and writing it in matrix form, we have $\Gamma(t) = \lim_{\Delta t \downarrow 0} \frac{P(t, t+\Delta t) - P(t, t)}{\Delta t} = \lim_{\Delta t \downarrow 0} \frac{P(t, t+\Delta t) - I}{\Delta t}$. From this it follows that for small Δt , we have $P(t, t + \Delta t) = I + \Gamma(t)\Delta t$.

³For more formal treatment of product integral see Gill and Johansen (1990) and for a lucid exposition with some applications, see Gill (2005).

$\lambda_{hj}^0(t)$ is known as the *baseline hazard function*. The specification of transition intensity in Eq. (6) is known as the *proportional hazard model*. It aggregates the effects the regressors linearly as a measure of some kind of latent factor, and that latent factor shifts the baseline hazard proportionately, i.e., the effect on hazard is uniform over time. Two papers — Fleming (1978) and Aalen and Johansen (1978) — independently extended the Kaplan-Meier nonparametric product limit estimator from survival analysis to the multi-state time to event models. While Fleming (1978) gave the estimator for complete data, Aalen and Johansen (1978) gave the estimator for censored data. To describe the Aalen-Johansen estimator, let me introduce some concepts and notation. For each individual $i, i = 1, 2, \dots, n$ and corresponding to each transient health state, $h, h = 1, 2$, define two types of stochastic processes: (1) the counting processes $N_{hj,i}(t)$ denoting the **observed** number of transitions from health state h to health state j that the individual i has made by time t —which in our case is either 0 or 1, since by assumption when an individual exits a health state, the individual does not return to it in future ; and (2) $Y_{h,i}(t)$, taking value 1 if individual i is at risk at time t for transition to another possible health state, and taking value 0 otherwise.

Let us focus on one transition $h \rightarrow j$. Denote by $\bar{N}_{hj}(t) = \sum_i^n N_{hj,i}(t)$, a counting process measuring the number of transitions of the $h \rightarrow j$ in the sample at time t , $\bar{Y}_h(t) = \sum_i^n Y_{h,i}(t)$, a counting process measuring the number of individuals in the sample at risk for a transition at time t , and $\bar{M}_{hj}(t) = \sum_i^n M_{hj,i}(t)$. In any empirical study the data will be at the discrete times, say in ordered times $0 = t_0 < t_1 < \dots < t_m$. At each time t_i , we calculate

$$\hat{\lambda}_{hj}(t_i) = \frac{\Delta \bar{N}_{hj}(t_i)}{\bar{Y}_h(t_i)}, j \neq h, \quad (7)$$

Without covariates, the *Nelson-Aalen non-parametric estimate* of the integrated intensity functions is given by, for each $h = 1, 2$

$$\begin{aligned} \hat{\Lambda}_{hj}(t) &= \sum_{i:t_i \leq t} \hat{\lambda}_{hj}(t_i), j \neq h, \\ \hat{\Lambda}_{hh}(t) &= - \sum \hat{\Lambda}_{hj}(t) \\ \hat{\Lambda}_{hj}(t) &= 0 \text{ for all other } h, j \text{ combinations} \end{aligned}$$

The *Aalen-Johansen estimator* $\hat{P}(s, t), s, t, \in T, s < t$ for the transition probabilities is obtained by substituting for each component hj the Nelson-Aalen estimates $\hat{\Lambda}_{hj}(t)$ and then applying the product integral formula Eq. (5) as follows

$$\hat{P}(s, t) = \prod_{s < u < t} (I + d\hat{\Lambda}(u)) = \prod_{i: t_i \leq t} (I + [\hat{\Lambda}(t_i) - \hat{\Lambda}(t_{i-1})]). \quad (8)$$

With covariates one obtains the Cox partial likelihood estimate for $\hat{\beta}_{hj}$ for each transition $h \rightarrow j$ separately and then computes an weighted risk set defined by

$$\bar{Y}_{hj}^*(t) = \sum_{i=1}^n Y_{hj,i}(t) \exp(\hat{\beta}_{hj}' X_{h,i}^0). \quad (9)$$

The estimates of cumulative intensities with covariates are obtained from Eq. (7) by replacing, $\bar{Y}_h(t)$ with $\bar{Y}_{hj}^*(t)$.

Nelson-Aalen estimator has nice statistical property. For instance, using Martingale calculus, it can be shown that the estimator is asymptotically unbiased. Using the results from Martingale theory, one can derive the formula for variance-covariance estimates of parameter estimates and the normalized estimate is normally distributed (central limit theorem holds for normalized parameter estimates), see for details, (Aalen, Ø. Borgan, et al., 2008; Andersen, O. Borgan, et al., 1993; Fleming and Harrington, 2005).

The likelihood of the sample (for details, see (Andersen, O. Borgan, et al., 1993; Andersen and Perme, 2008; Commenges, 2002)),

$$L(\theta) = \prod_i \prod_{h=1,2; j=2,3,4; h \neq j} \left(\prod_t \lambda_{hj,i}(t X_{h,i})^{\Delta N_{hj,i}(t)} \right) \exp\left(-\int_0^{T_{h,i}^*} \lambda_{hj,i}(u X_{h,i}) du\right) \quad (10)$$

I will use this likelihood function in construction of the log-likelihood loss function for the neural network models.

I use the R package, *mstate*, developed and described by the authors in Wreede et al. (2010) for the estimation of the parameters and their standard errors?

The Aalen-Johnson nonparametric estimates without any covariates of the transition probabilities, are shown in Table 2 and plotted in Figure 2.

For the Cox regression parameter estimates, I have used both the R package *mstate* (see, Wreede et al. (2010) for details) and also used the SAS procedure *phreg* (both produced the same estimates)

and used the *mstate* package to estimate all the transition probabilities (SAS does not have readily available procedure for this purpose). The parameter estimates are shown in [Table 3](#) and [Table 4](#).

3 The data set and the variables

I use the Health and Retirement Study (HRS) dataset for empirical analysis. A lot has been reported on the family of HRS datasets—about its structure, purpose, and various modules collecting data on genetics, biomarkers, cognitive functioning, and more, see for instance ([Juster and Suzman, 1995](#); [Sonnega et al., 2014](#); [Fisher and Ryan, 2017](#)). The first survey was conducted in 1992 on a representative sample of individuals living in households i.e., in non-institutionalized, community dwelling, in the United States from the population of cohort born during 1931 to 1941 and their spouses of any age. “The sample was drawn at the household financial unit level using a multistage, national area-clustered probability sample frame. An oversample of Blacks, Hispanics (primarily Mexican Americans), and Florida residents was drawn to increase the sample size of Blacks and Hispanics as well as those who reside in the state of Florida’’, [Fisher and Ryan \(2017\)](#). The number of respondents were 13,593. Since 1992, the survey were repeated every two years, each is referred to as a wave of survey. New cohorts were added in 1993, 1998, 2004 and 2010, ending the survey up with the sample size of 37,495 from around 23,000 households in wave 12 in 2014. RAND created many variables from the original HRS data for ease of use. I create my dataset and all the variables with a few exceptions mentioned below from the RAND HRS dataset version P. The details of the Rand HRS version P can be found in [Bugliari et al. \(2016\)](#).

As mentioned in the introduction, I define the disability health state to be one that qualifies one to be on the disability programs OASDI or SSI. The data on disability is self-reported. Later I plan to use the Social Security Administration’s matched administrative data on this variable and earnings variables not included here. The matched data will, however, reduce the sample size to half, as only 50 percent of the respondents are used for matching HRS with SSA Administrative data. The HRS data collected information on if and when the doctor diagnosed that the respondent has any of the severe diseases such as high blood pressure, diabetes, cancer, lung disease, heart attack, stroke, psychiatric disorder and severe arthritis. I drop respondents who received disability before the first survey year 1992 and I also drop the spouses in the sample who were not born between 1931 to 1941, that is the respondents in our sample are between age 51 to 61 and not disabled or dead in 1992. I ended up with the final sample size of 9,493 for this analysis.

[Table 1](#) provides a few characteristics of the data.

Table 1: Summary of the health status of the individuals in the sample over the survey years.

Survey Year	Alive:normal health	Alive:diseased	Became disabled	Died before disability	65+: censored	Total
1992	3026	6483	92	0	0	9601
1994	2591	6608	166	144	0	9509
1996	2291	6623	139	146	0	9199
1998	1726	5478	134	123	727	8188
2000	1279	4313	86	113	741	6532
2002	848	3148	54	58	759	4867
2004	468	1893	34	48	781	3224
2006	132	636	4	14	795	1581

3.1 Variables

The demographic variables **White** and **Female** have the standard definition. The variable **College** is a binary variable taking value 1 if the respondent has education level of college and above (does not include some college), i.e., has a college degree and more and taking value 0 otherwise.

cesd: I used a score on the Center for Epidemiologic Studies Depression (CESD) measure in various waves that is created by RAND release of the HRS data. RAND creates the score as the sum of five negative indicators minus two positive indicators. “The negative indicators measure whether the Respondent experienced the following sentiments all or most of the time: depression, everything is an effort, sleep is restless, felt alone, felt sad, and could not get going. The positive indicators measure whether the Respondent felt happy and enjoyed life, all or most of the time.” I standardize this score by subtracting 4 and dividing 8 to the RAND measure. The wave 1 had different set of questions so it was not reported in RAND HRS. I imputed it to be the first non-missing future CESD score. In the paper, I refer the variable as *cesd*. [Steffick \(2000\)](#) discusses its validity as a measure of stress and depression.

cogtot: This variable is a measure of cognitive functioning. RAND combined the original HRS scores on cognitive function measure which includes “immediate and delayed word recall, the serial 7s test, counting backwards, naming tasks (e.g., date-naming), and vocabulary questions”. Three of the original HRS cognition summary indices—two indices of scores on 20 and 40 words recall and third is score on the mental status index which is sum of scores “from counting, naming, and vocabulary tasks” — are added together to create this variable. Again due to non-compatibility with the rest of the waves, the score in the first wave was not reported in the RAND HRS. I have imputed it by taking the first future non-missing value of this variable.

bmi: The variable body-mass-index (BMI) is the standard measure used in the medical field and HRS collected data on this for all individuals. If it is missing in 1992, I impute it with the first future non-missing value for the variable.

behav_smoke: This variable is constructed to be a binary variable taking value 1 if the respondent has reported yes to ever smoked question during any of the waves as reported in the RAND HRS data and then repeated the value for all the years.

behav_vigex: The RAND HRS has data on whether the respondent did vigorous exercise three or more days per week. I created this variable in each time period to take the value 1 if the respondent did vigorous exercise three or more days per week.

3.2 Statistical Results

Table 2: Estimated transition probabilities for transitions ioj by duration of stay in state 1 and 2.

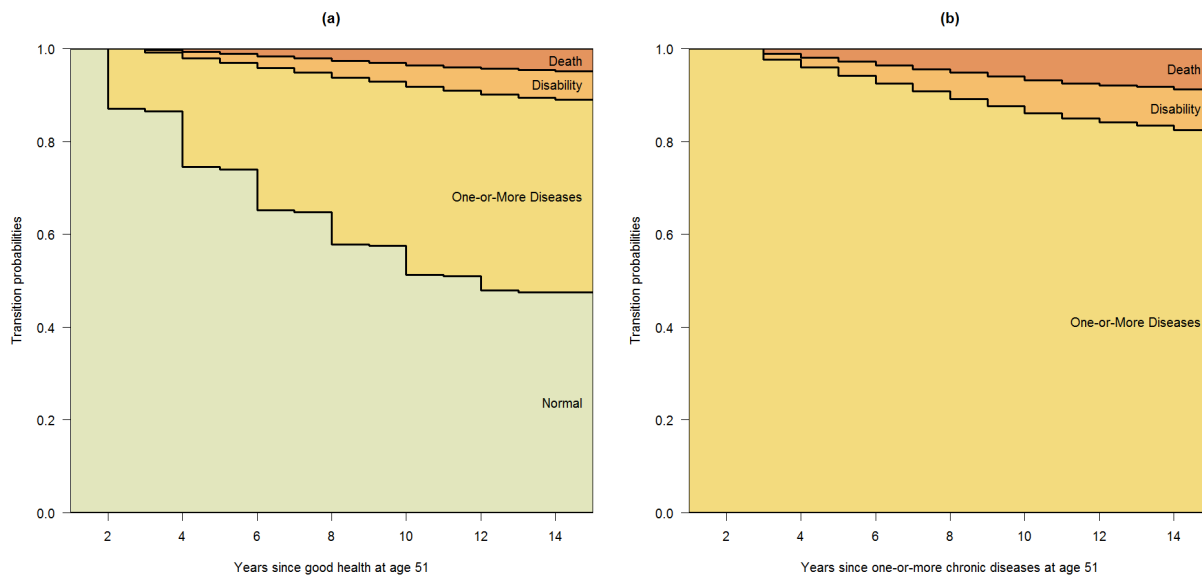
Duration	1 \rightarrow 1	2 \rightarrow 2	1 \rightarrow 2	1 \rightarrow 3	2 \rightarrow 3	1 \rightarrow 4	2 \rightarrow 4
0	1.0000	1.0000	0.0000	0.0000	0.0000	0.0000	0.0000
1	0.8711	1.0000	0.1289	0.0000	0.0000	0.0000	0.0000
2	0.7543	0.9767	0.2367	0.0062	0.0127	0.0027	0.0106
3	0.6570	0.9430	0.3105	0.0201	0.0292	0.0123	0.0279
4	0.5786	0.9082	0.3651	0.0331	0.0477	0.0232	0.0441
5	0.5126	0.8730	0.4087	0.0446	0.0659	0.0341	0.0610
6	0.4761	0.8421	0.4212	0.0562	0.0792	0.0465	0.0787
7	0.4667	0.8229	0.4116	0.0681	0.0896	0.0536	0.0874

With only demographic covariates (most that can be done with the Administrative data) the parameter estimates in [Table 3](#) show that significantly lower risks of transitions 1 \rightarrow 4; 2 \rightarrow 3; 2 \rightarrow 4 for whites and 1 \rightarrow 4 and 2 \rightarrow 4 for women. This may entail that the genetic make-up of being white or female sex yield favorable genetic predisposition to have better health outcomes and longer life. This is a misleading inference as we will see next that when we control for epigenetic factors that the biomedical literature pointed out to have significant effects on aging process and health outcomes, the above effects disappear.

[Table 4](#) shows the Cox regression coefficient estimates of the effects of various factors on the risk of having transitions $h \rightarrow j$ from health status $h = 1, 2$ to health status $j, j = 2, 3, 4$.

These estimates show that the parameter estimates for the demographic covariates in [Table 3](#) are biased as they are capturing the effects of excluded epigenetic and behavioral factors in that model.

Figure 2: Transition probabilities (a) from normal health state, (b) from one-or-more diseased health state



After controlling for these epigenetic and behavioral factors, the significance of those effects disappear. Furthermore, women show significantly lower probability of transition from diseased health states onto the disability health state.

Most important factors in Table 4 are cesd, measuring depression and stress and college graduation or higher level of education, with positive effect on all transitions with the exception of no effect on transition from normal health to death.

Other important factors are smoking, with significant adverse effect on transitions, and exercising three or more times regularly has significant favorable effect on most transitions.

Table 3: Estimates of Cox regression models separately for each transition with demographic variables only

	1->2	1->3	1->4	2->3	2->4
White	-0.028 (0.067)	-0.328 (0.230)	-0.716 ** (0.265)	-0.387 *** (0.102)	-0.426 *** (0.102)
Female	0.052 (0.050)	-0.216 (0.190)	-0.522 * (0.251)	-0.029 (0.092)	-0.493 *** (0.093)
#obs	3583	3695	3652	6856	6855
#events	1602	112	69	476	475
R squared	0.000	0.001	0.003	0.002	0.006
logLik	-12671.44	-854.58	-518.80	-4038.09	-3994.91
AIC	25346.88	1713.15	1041.60	8080.17	7993.81

*** p < 0.001; ** p < 0.01; * p < 0.05.

Table 4: Estimates of Cox regression models separately for each transition with health measures

	1->2	1->3	1->4	2->3	2->4
cesd	0.561 *** (0.120)	1.980 *** (0.352)	0.126 (1.109)	1.295 *** (0.155)	0.547 * (0.256)
bmi	0.042 *** (0.006)	-0.005 (0.025)	0.012 (0.050)	0.025 ** (0.008)	-0.021 (0.014)
cogtot	-0.003 (0.006)	-0.067 ** (0.022)	0.015 (0.048)	-0.035 *** (0.010)	-0.009 (0.015)
behav_smoke	0.045 (0.054)	0.258 (0.226)	2.511 * (1.028)	0.381 *** (0.108)	0.817 *** (0.180)
behav_vigex	-0.197 ** (0.075)	-1.000 *** (0.235)	-1.169 * (0.485)	-0.610 *** (0.102)	-1.199 *** (0.145)
White	0.045 (0.072)	0.098 (0.269)	-0.494 (0.498)	-0.117 (0.112)	-0.325 * (0.165)
College	-0.111 (0.069)	-0.887 * (0.406)	-1.114 (0.760)	-0.684 *** (0.195)	-0.744 ** (0.262)
Female	0.085 (0.055)	-0.273 (0.221)	-0.686 (0.470)	-0.080 (0.101)	-0.252 (0.146)
#obs	3239	3334	3262	6165	5926
#events	1500	95	23	446	207
R squared	0.029	0.025	0.008	0.037	0.022
logLik	-11647.65	-671.06	-151.05	-3626.20	-1626.01
AIC	23311.31	1358.12	318.10	7268.41	3268.01

*** p < 0.001; ** p < 0.01; * p < 0.05.

4 Long-short memory (LSTM) recurrent neural network (RNN) model of health outcomes process

I first give an overview of the neural network framework. I then briefly review a selected few papers that used the neural network framework to estimate and predict the time-to-event probabilities in a single event set-up. The main focus of the papers, including this paper, is to relax the proportionality and linearity assumption of the statistical survival models, see Eq. (6) for the proportionality assumption. I then explain why a long-short memory (LSTM) recurrent neural network (RNN) framework is more appropriate for predictive models of time-to-event probabilities in sequential set-up. I describe the *c-index* criterion that I use to compare the performance of various models and compare the performance of the multistate statistical model and a feed-forward neural network model of competing risk model by Lee et al. (2018) with the LSTM-RNN model of this paper. I code these three models using Keras and Tensorflow 2.0 modules in Python and estimate them using the Health and Retirement Survey data mentioned above. In the following section, I report the findings.

4.1 The basics of neural network

I first give an overview of the neural network framework. I then briefly review a selected few papers that used the neural network framework to estimate and predict the time-to-event probabilities in a single event set-up. The main focus of the papers, including this paper, is to relax the proportionality and linearity assumption of the statistical survival models, see Eq. (6) for the proportionality assumption. I then explain why a long-short memory (LSTM) recurrent neural network (RNN) framework is more appropriate for predictive models of time-to-event probabilities in sequential set-up. I describe the *c-index* criterion that I use to compare the performance of various models and compare the performance of the multistate statistical model and a feed-forward neural network model of competing risk model by Lee et al. (2018) with the LSTM-RNN model of this paper. I code these three models using Keras and Tensorflow 2.0 modules in Python and estimate them using the Health and Retirement Survey data mentioned above. In the following section, I report the findings.

Neural network is a highly parameterized universal function approximator of the form $\hat{y} = f(x; w)$, x is a set of inputs, and w is a vector of parameters. This is of the same nature as a statistical model. More precisely, suppose we have data on a set of individuals of the type (x, y) , where x is a vector of individual characteristics, and y is a vector of output levels and w is a set of parameters. The output could be a categorical variable for classification problems, it could be a probability distribution

over finite classes, as in our case, or it could be a continuous variable for regression problems. The data generating process for y as a function of x , is not known. The goal is to approximate that unknown data generating function. This is the problem that both statistics and neural network find computational solutions to. In neural network, the problem is to design a *neural network architecture* of the approximating function $\hat{y} = f(x, w)$ and find a suitable *learning algorithm* to learn the parameter values w of the network using a training set of examples. This trained network can then be used to predict y for an individual given his characteristics x .

The popularity and wide applicability of neural network lies in the fact that it designs the approximator in a hierarchy of functions, joined together by compositions of functions, that renders good properties in terms of ease of computation and closeness of function approximation. Most neural network models have the following type of hierarchical functional form:

$$\hat{y} = f(x; w) \equiv f_{w^L}^L \circ \dots \circ f_{w^1}^1(x). \quad (11)$$

Each function corresponds to a layer of artificial neurons. The role of each neuron is to perform simple calculations and then pass on that to the next layer of neurons.

Neurons in each layer get signals which are the outputs of the neurons of the previous layer (also known as activation levels) that it is connected with. It sums them, I denote this sum with z and apply an activation function to produce an output also known as *activation level* which I denote by a . The activation level a will then be passed on as an input to a neuron that it is connected to in the next layer. The neurons of the last layer will compute the output level taking the activation levels of the connected neurons of the previous layer.

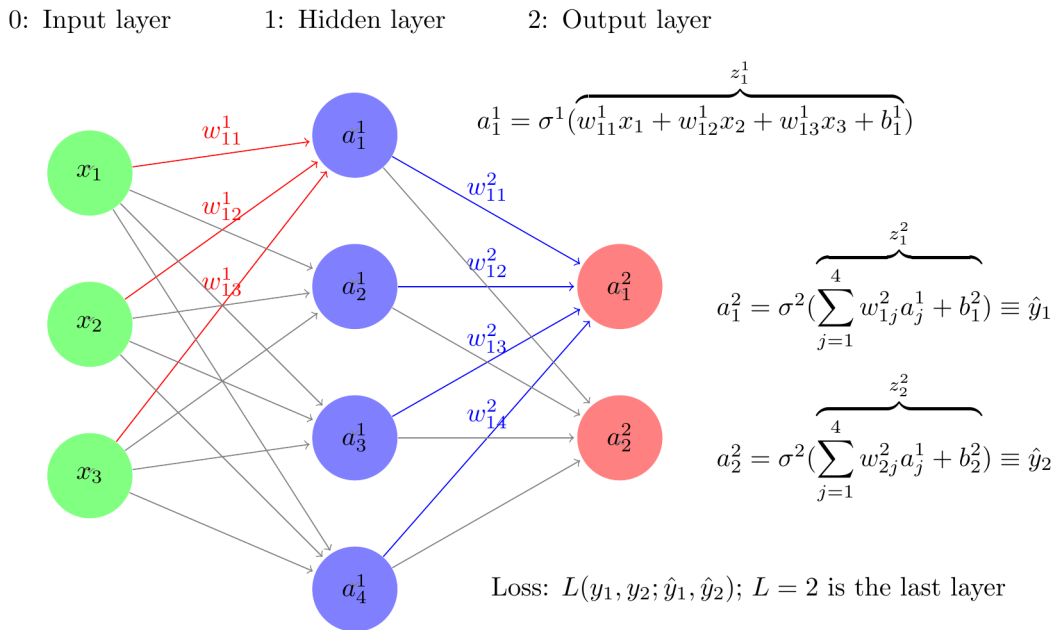
Consider a simple neural network architecture depicted in [Figure 3](#). It has three layers — layer 0: input layer, layer 1: hidden layer, and layer 2: output layer. Last layer in the text is denoted by L , and hence $L = 2$. Layer 0 has three input neurons. The second layer has 4 neurons. and last layer has two neurons corresponding to the two output levels, in our case probability of two events. In this neural network, the hierarchical function specification is of the form:

$$f(x; w) = \sigma^2(z^2(\sigma^1(z^1(x, w^1)), w^2)) \equiv f_{w^2}^2 \circ f_{w^1}^1(x). \quad (12)$$

The function $z^i(a^i, w^i) = w^i \cdot a^{i-1}$ at each layer i is a linear aggregator. In the notation, z^i is a vector of functions, each component of which corresponds to a neuron of the i th layer. The function σ^i is a squashing function of the same dimension as z^i , each component having the same function real

valued function of one variable, known as *activation function*, which squashes the value of z^i to a range such as $(0, 1)$ for a sigmoid activation function, or $(-1, 1)$ for the *tanh* activation function. The value of the activation function, $a^i = \sigma^i(z^i)$ is known as the *activation level* of the neurons of the i th layer. The activation levels of the 0th layer, $a^0 = x$, the inputs, fed to the neural network from outside. The operation on the right is also performed component-wise for each neuron at the i th layer it computes the weighted sum of the activation levels (outputs) of the neurons of the previous layer that the neuron of the i th layer is connected to. The weights used are specific to the neuron of the i th layer. An activation function σ^i which generally taken to be same for all the neurons of the i th layer) is applied to this aggregated value z^i . These activation functions do not have any unknown parameters that need to be estimated. These two computations—aggregation and activation—are shown as a composite mapping $f_{w^i}^i$ for the neurons of the i th layer.

Figure 3: MLP architecture.



There are different types of neural network, depending on the functional form of f in Eq. (11). A *deep feed forward neural network*, also known as a *feedforward neural network with hidden layers* or as a *multilayer perceptron (MLP)* is a network architecture in which there is no feedback of any neurons to itself or to others in the same layer. The activation levels of the neurons only feed forward to the neurons in the next layer. This is the reason also why these type of neural networks are called feed forward, as opposed to the recurrent neural network consider later that allow feedback. The MLP has good computational properties and an MLP is a great universal functional approximator: It is shown [Hornik et al. \(1989\)](#) that with a sufficient number of layers

in a hidden layer can approximate a function to any level of precision desired. So a MLP can be used to approximate the true data generating process as closely as one wants. How does one find such a network, i.e., how does one choose the weights of the network.

To get a good approximation, the artificial neural network contains hundreds of thousands of deep parameters w , how does one train the network, i.e., how to learn the the parameter values. The learning is done by choosing the weights to minimize a loss function together with a nonnegative regularization term (in statistical term a regularization term corresponds to shrinkage estimator).

$$L(y, f(x, w)) + \lambda C(w). \tag{13}$$

In the present context of learning about a probability distribution of time-to-events, an appropriate loss function is to take the negative log-likelihood of the sample and the additive regularization term $C(w)$ to be $\|w\|_2$. In this case the loss function with regularization term has a Bayesian interpretation—it is the posterior log-likelihood with normal prior distribution for the parameters w . The choice of w to minimize the loss is done by a gradient descent method. The neural network architecture Eq. (11) yields a very convenient fast and automatic computation of the gradients $\partial L/\partial w$ using an algorithm known as back-propagation algorithm, used pretty much in all types of neural networks. Two steps in this algorithm are assume an initial value of the weights w , first compute all the activation levels starting at the input layer, i.e., $f_w^1(x)$ forward through the layers 2, 3, ..., L in Eq. (11), i.e., go through layer superscripts forward. In the next step, compute the gradients of the weight parameters w of various layers, starting from the last layer, backward in layers, i.e., decreasing order in the subscripts in Eq. (11). Once all gradients are computed, weights are adjusted using a steepest descent algorithm. Further details are omitted, since I do not implement any type of back propagation algorithm in this paper.

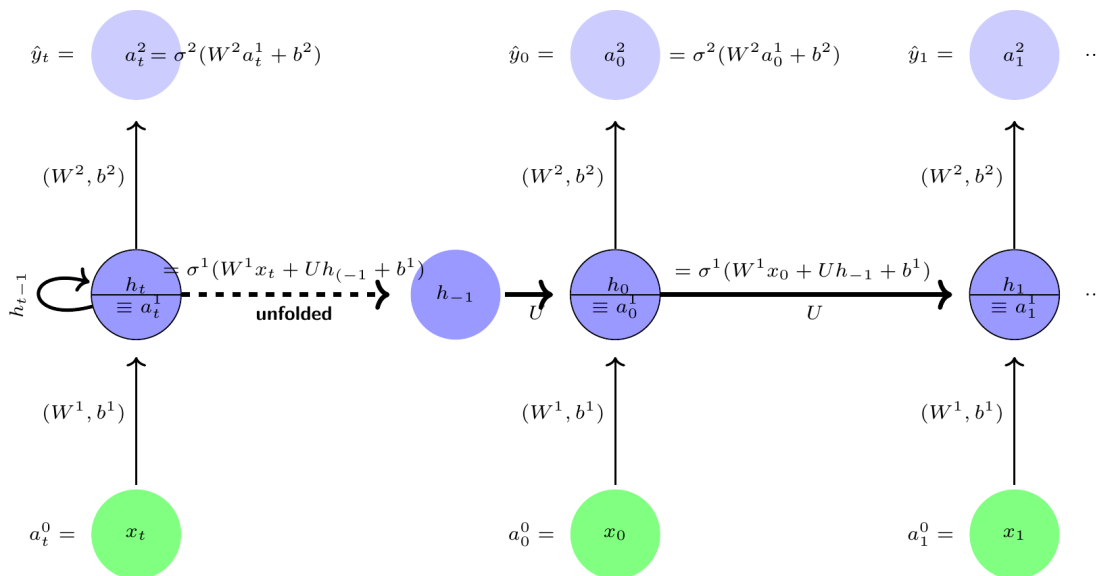
4.2 Previous neural network models of survival analysis

The main objective of the neural network approach to survival analysis is the relaxation of restrictive proportional hazard assumption in Eq. (6). The following papers use a MLP architecture to estimate the survival probabilities with covariates fixed at the beginning of the study period. (Faraggi and Simon, 1995; Biganzoli et al., 1998; Fernandez et al., 2016; Katzman et al., 2018; Lee et al., 2018; Ranganath et al., 2016; Ren et al., 2019; Zhao et al., 2019). Except for Lee et al. (2018) who consider a competing risk model, all other papers consider two-state, alive-death, models.

Two main limitations of the previous literature is that they use feed-forward neural network architecture to specify the functional approximator for the hazard function in Eq. (6). These are static model in the sense that the input vectors of an individual are given at the begin of the period and the neural network model predict the probabilities of the event for a fixed number of periods. This is problematic because prediction of probabilities at a given period use all the information from the future periods. It cannot handle training using data in which some individuals do not have information for all periods. More importantly, these models cannot incorporate new information that may come after training the model to update the time-to-event prediction probabilities. Another shortcoming of these models is that they are restricted to analyzing life histories of only one transition between two health states—from alive to death or at the most to various causes of death in the language of survival analysis. As I have mentioned in the introduction, to better predict the probabilities of time to certain health come of interest, it is important to analyze dynamic paths through other intermediary health outcomes and the factors that affect the dynamics of those paths. I argue that a long-short memory recurrent neural network is more appropriate model to that end. I explain this framework next.

4.3 Formulation within Recurrent Neural Network with LSTM

Figure 4: Showing a simple recurrent neural, with layer 0 is the input layer, layer 1 is the recurrent hidden layer, layer 2 is another pure hidden layer, layer 3 is the output layer, for each time point (t), layers are denoted as superscript, and the subscript in parenthesis is time.



Let x_t be a set of measurements, or characteristics and y_t , a set of health outcomes are measured

repeatedly over time $t, t = 0, 1, \dots, T$. A recurrent neural network (RNN) allows a feedback link among neurons in a hidden layer. A typical RNN is shown in the left panel and the unfolded version in the right panel of [Figure 4](#). This neural network is a simple three layer architecture similar to MLP architecture shown earlier, with one main difference is that the hidden layer has a feedback connection linked to itself, i.e., a cyclical or recurrent connection and there is weight sharing, i.e., the weights of all the layers for a given time period t is constant for all time periods. The recurrent connections in the RNN architecture effectively creates an internal memory of the effects of previous inputs that can affect the current and future outcomes. It is also interesting to note that the functional forms of the RNN are similar to Kalman filtering models, but more general in some respects and restrictive in some other respects.

The hidden layer 1: The activation levels of the hidden recurrent layer at time-step t is produced by combining activation levels h_{t-1} of the previous period's hidden recurrent layer and input levels x_t and then applying the activation function σ^1 as follows: For $t = 0, 1, 2, \dots, T$ Layer 1 (hidden recurrent layer):

$$\begin{aligned}
 h_t &= \sigma^1(W^1 \cdot x_t + U \cdot h_{t-1} + b^1) \\
 a_t^1 &\equiv h_t \text{ (for notational convenience)} \\
 h_{-1} &= \text{a user supplied initial activation level}
 \end{aligned}$$

Layer 2 (Output):

$$\begin{aligned}
 a_t^2 &= \sigma^2(W^2 \cdot a_t^1 + b^2) \\
 \hat{y}_t &\equiv a_t^2 \text{ (notational convenience)}
 \end{aligned}$$

The pseudo computational graph of the network is shown in [Figure 4](#).

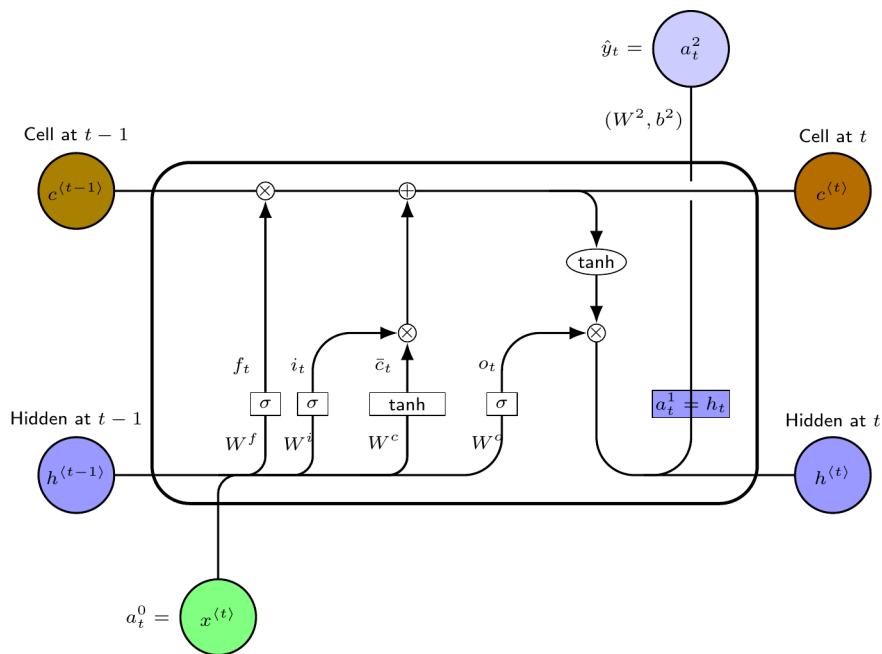
It may seem a minor difference between the architectures of an MLP and an RNN, but difference in their capabilities are very important and crucial for sequential learning. For instance, an RNN can map the entire history of previous inputs to each output as compared to an MLP which can map only current input to output vectors. Similar to the result that an MLP is a universal function approximator mentioned earlier, with a sufficiently large number of hidden self-connected neurons, an RNN can approximate extremely closely any sequence-to-sequence mapping, see ([Graves et al., 2014](#); [Hammer, 2000](#); [Siegelmann and Sontag, 1992](#)). There are various types of RNN, for details see ([Graves, 2012](#); [Lipton et al., 2015](#)). In many situations, especially in drug-discovery, the study samples may have small size, an important extension of an RNN framework that handles this problem is ([Altae-Tran, 2016](#); [Altae-Tran et al., 2017](#)).

The computation of the gradients is done using a modification of the back propagation method known as back propagation through time, which involves multiplication of gradients less than one many times over the time steps tending to vanishing gradient problem, meaning the gradient of the weights representing memories of the far back in time tend to zero or vanish. A similar problem arises when the gradients larger than 1 are multiplied many times leading to exploding gradient problem.

Quite a few other fixes are proposed for the vanishing and exploding gradient problems of RNN models, such clipping of gradients, and various other types of extensions such as [Elman \(1990\)](#) are proposed. But the long-short memory extension of the RNN framework by [Hochreiter and Schmidhuber \(1997\)](#) proved to be very useful and successful in many applications which I use for this paper.

4.4 LSTM memory cell

Figure 5: Showing a typical memory cell at time-step t



The main innovation in this approach is to replace the hidden recurrent cells in an RNN with a memory cell with three gates—an input gate, an output gate and a forget gate—and a recurrent cell memory together with the original hidden layer memory cell. The functional form of a typical LSTM memory cell is as follows

Layer 1 (Hidden layer with recurrent neurons):

$$f_t = \sigma(W^f \cdot x_t + U^f \cdot h_{t-1} + b^f)$$

$$i_t = \sigma(W^i \cdot x_t + U^i \cdot h_{t-1} + b^i)$$

$$o_t = \sigma(W^o \cdot x_t + U^o \cdot h_{t-1} + b^o)$$

$$\bar{c}_t = \tanh(W^c \cdot x_t + U^c \cdot h_{t-1} + b^c)$$

$$c_t = f_t * c_{t-1} + i_t * \bar{c}_t$$

$$h_t = o_t * \tanh(c_t)$$

$$a_t^1 \equiv h_t \text{ (for notational convenience)}$$

$$h_{-1}, c_{-1} = \text{user supplied initial activation levels}$$

Layer 2 (Output): This is same as the RNN output layer.

$$a_t^2 = \sigma^2(W^2 \cdot a_t^1 + b^1)$$

$$\hat{y}_t \equiv a_t^2 \text{ (notational convenience)}$$

The pseudo computational graph of the network is shown in [Figure 5](#).

5 Prediction and performance of the model

As a multistate time to event model has censored data, the standard metrics such as R^2 , root mean square error are not applicable for performance measurement. [Harrell et al. \(1982\)](#) introduce the measure of concordance index, known as c-index, extending the concept of the area under the ROC curve, which was further refined by [Antolini et al. \(2005\)](#) to be applicable for survival models with time-varying covariates. Most machine learning models use this criterion to compare the performance of various neural network models of survival analysis.

To get the basic idea behind the c-index, suppose two individual in the dataset are in a particular health state, both of whom are at risk for exits to another possible health state. Suppose the first individual exits at time t and the second individual did not exit at time t or earlier. The first individual's estimated survival probability at time t should be smaller than the second individual who did not exit. A good model should have this maintained for most or all such comparable pairs. The c-index measures the proportion people for whom this is true out of all the people who could be compared.

I use the c-index criterion to discriminate the performance of various models in predicting time-to-event probabilities. The c-index estimates for three three models are shown in the last row

of Table 5. Judging from the c-index estimates, it appears that neural network models do better job in prediction of time to event probabilities than Aalen-Johansen estimators of the multistate statistical model. Furthermore, the LSTM-RNN model of this paper does perform slightly better.

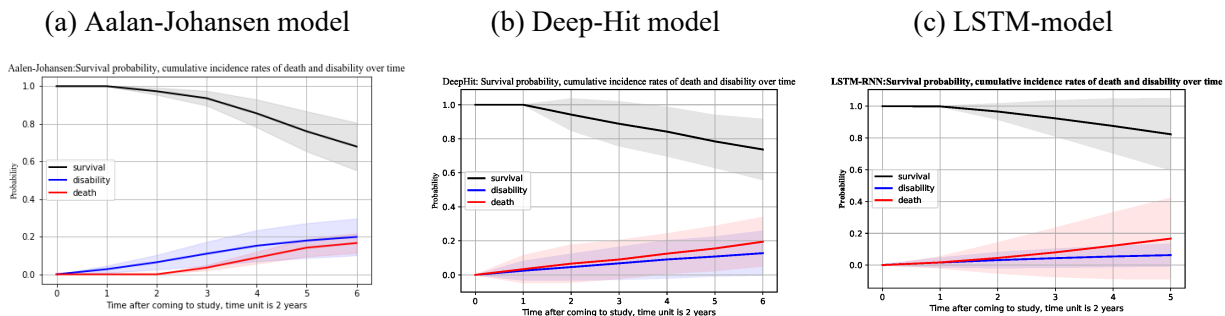
Table 5: Average of the Predicted cumulative incidence rates of disability and death for three methods for the test sample.

T	Statistical multistate model		Lee et al Deephit model		LSTM-RNN model	
	Disability	Death	Disability	Death	Disability	Death
0.000	0.000	0.000	0.000	0.000	0.001	0.000
1.000	0.027	0.000	0.035	0.000	0.017	0.017
2.000	0.064	0.000	0.052	0.004	0.032	0.045
3.000	0.109	0.035	0.120	0.026	0.044	0.080
4.000	0.152	0.088	0.247	0.050	0.055	0.122
5.000	0.179	0.141	0.271	0.133	0.063	0.167
6.000	0.199	0.166	0.287	0.155	0.070	0.208
7.000	0.206	0.197	0.496	0.504	0.077	0.248
c-index	0.476		0.748		0.756	

I also show in Table 5 the average of the predicted time to event probabilities of the individuals in the test data set, and in Figure 6, I plotted the average probabilities together with the confidence intervals. Looking at the variability in the graph, and direct comparison of probabilities for the individual cases, it appears that neural network models might be over-fitting the data, in spite of using the L_2 regularization. More investigation is needed.

So I based my comparison of the models using the c-index.

Figure 6: Means and mean \pm standard deviation intervals of the predicted cumulative incidence rates of disability and death in the test sample



6 Conclusion

Individuals follow different health trajectories over their lives. Many factors along a health trajectory up to a given time period affect their health state at the time and the future paths that the individual will follow. Such sequential health developments conditioned by the health related behaviors along the path determine the risks of diseases, disability and death. This paper studies statistical and artificial neural network models for predicting time to an event of interest—disability—with death as a competing risk event during the mid ages and how different factors along the health trajectory affect these probabilities. The paper uses the Health and Retirement Study data for estimation of the models. The paper discusses the merits and demerits of these two approaches and compare their performances.

The paper formulates the life histories of individuals as a stochastic process through a finite set of health states. The main problem is to estimate the transition probabilities between health states and the effects of various time varying covariates on transition probabilities. The effect of covariates is parameterized through “transition intensity functions” in both statistical and neural network models. Statistical multistate models assume that in each transition, the covariates linearly affect a latent health measure and that latent health measure shifts a baseline hazard function (or transition intensity function) proportionately up and down by a multiplicative factor. This type of specification of transition intensities is known as Cox proportional hazard model. In many situations, including the present, this proportionality assumption imposes strong structure and could be far from the structure of the true data generating process. The multi layer perceptron (MLP), a kind of deep neural network (explained in the paper), is known to be an excellent ‘universal function approximator’. Utilizing this insight, recently a few papers use MLP models to relax the proportionality assumption. This paper argues that the prediction of progression through health states and estimation of the probabilities of health events of interest such as disability and death are better done in a recurrent neural network (RNN) model with long-short memory cells to capture the effects of long lagged health related behaviors and outcomes that affect the current and future health outcomes. The paper compares the performances of a statistical multistate model, and an MLP based competing risk model of [Lee et al. \(2018\)](#) and the LSTM-RNN model proposed in this paper. The paper finds that performance, measured by c-index, is much better for both the neural network models and is slightly better for the LSTM-RNN model. Since an LSTM-RNN neural model is more suitable for prediction in a sequential set-up, the findings prescribe that it is better to use an LSTM-RNN type of neural network model to estimate and predict transition probabilities than an MLP based neural network model.

For all three models, the paper included some epigenetic factors (that include health related behaviors), demographic factors, education level, some biomarkers like BMI, CESD and cognition and depression and stress level—time varying for the LSTM-RNN and multi-state statistical model and time fixed at the first period for the Lee et al. (2018) model. At present there are no neural network software package that can estimate and compare the relative importance of the covariates on transition probabilities. The statistical models are good at that. More research in neural network along this line will be useful.

From the estimates of the statistical model, the paper finds that college graduates have significantly lower probability of all the transitions. The variable CESD measuring the level of depression and stress has significant positive effects on transiting from normal health to diseased health state, from normal health to becoming disabled and from diseased health state to become disabled or to death. The other most significant behavioral variables are smoking and sufficiently vigorous level of regular exercising. The smoking has significantly adverse effects and exercising has favorable effects on most transitions.

For unfamiliar readers, the paper gives introduction to estimation of multistate stochastic models and the MLP and LSTM-RNN type of neural networks models when the data has censored observations.

While a statistical model is good for studying the effects of various covariates on time to event probabilities, judging from the performance based on the c-index, its performance as a predictive model is much worse than the two neural network models considered. Relaxation of the proportionality assumption to include more general specifications is a useful direction. For neural network models, it will be very useful to study the relative strength and statistical significance of various covariates on the outcomes such as transition probabilities through different health states that the models predict.

Appendix

MLP learning algorithm

The sets of neurons at layer $\ell = 0, 1, 2$ are $N^0 = \{1, 2, 3\}$ corresponding the three inputs, $N^1 = \{1, 2, 3, 4\}$, i.e. 4 hidden units, and $N^2 = \{1, 2\}$ at the output layer corresponding to 2 outputs. Let $n^k = \#N^k$, i.e., the number of neurons in the k-th layer of the network, for $k = 0, 1, \dots, L$.

1. First compute all neuron outputs a_i^ℓ , $i \in N^\ell$ using the feedforward iteration: $a^\ell =$

$\sigma^\ell(W^\ell a^{\ell-1} + b^\ell)$, $\ell = 1, \dots, L$, with initial condition $a^0 \equiv x$ (input vector).

2. To calculate $\frac{\partial C}{\partial w_{ij}^\ell}$, and $\frac{\partial C}{\partial b_i^\ell}$, denote by $z^\ell \equiv W^\ell a^{\ell-1} + b^\ell$ and $\delta_i^\ell \equiv \frac{\partial C}{\partial z_i^\ell} \cdot \sigma^{\ell'}(z_i^\ell)$, and calculate first all δ_i^ℓ , $i \in N^\ell$, $\ell = 1, 2, \dots, L$, using the backpropagation iteration: $\delta^\ell = W^{\ell+1} \delta^{\ell+1} \odot \sigma^{\ell'}(z^\ell)$, $\ell = L - 1, L - 2, \dots, 1$, with the initial condition $\delta^L = \frac{\partial C}{\partial a^L} \odot \sigma^{L'}(z^L)$, where \odot denotes the Haddamard product, i.e., element wise product.
3. Now compute $\frac{\partial C}{\partial b_i^\ell} = \delta_i^\ell$ and $\frac{\partial C}{\partial w_{ij}^\ell} = \delta_i^\ell a_j^{\ell-1}$, $i = 1, 2, \dots, n^\ell$, $j = 1, 2, \dots, n^{\ell-1}$, $\ell = 1, 2, \dots, L$. In matrix notation, $\frac{\partial C}{\partial b^\ell} = \delta^\ell$ and $\frac{\partial C}{\partial W^\ell} = \delta^\ell a^{\ell-1}$, $\ell = 1, 2, \dots, L$.

Bibliography

- [1] Aalen, O. O., Borgan, Ø., et al. Survival and Event History Analysis. Springer New York, 2008. DOI: [10.1007/978-0-387-68560-1](https://doi.org/10.1007/978-0-387-68560-1) (cit. on p. 15).
- [2] Aalen, O. O. and Johansen, S. An Empirical Transition Matrix for Non-Homogeneous Markov Chains Based on Censored Observations, *Scandinavian Journal of Statistics*, **5**, no. 3 (1978), 141–150 (cit. on pp. 7, 13, 14).
- [3] Alisch, R. S. et al. Age-associated DNA methylation in pediatric populations, *Genome Research*, **22**, no. 4 (Feb. 2012), 623–632. DOI: [10.1101/gr.125187.111](https://doi.org/10.1101/gr.125187.111) (cit. on p. 5).
- [4] Altae-Tran, H. RNNs for Model Predictive Control in Unknown Dynamical Systems with Low Sampling Rates, *working paper, Stanford University* (2016) (cit. on pp. 8, 27).
- [5] Altae-Tran, H. et al. Low Data Drug Discovery with One-Shot Learning, *ACS Central Science*, **3**, no. 4 (2017), 283–293. DOI: [10.1021/acscentsci.6b00367](https://doi.org/10.1021/acscentsci.6b00367) (cit. on pp. 8, 27).
- [6] Andersen, P. K., Borgan, O., et al. Statistical Models Based on Counting Processes. Springer-Verlag, New York, 1993 (cit. on pp. 7, 15).
- [7] Andersen, P. K. and Perme, M. P. Inference for outcome probabilities in multi-state models, *Lifetime Data Analysis*, **14**, no. 4 (Sept. 2008), 405–431. DOI: [10.1007/s10985-008-9097-x](https://doi.org/10.1007/s10985-008-9097-x) (cit. on pp. 7, 15).
- [8] Antolini, L. et al. A time-dependent discrimination index for survival data, *Statistics in Medicine*, **24**, no. 24 (2005), 3927–3944. DOI: [10.1002/sim.2427](https://doi.org/10.1002/sim.2427) (cit. on p. 29).
- [9] Austad, S. N. and Fischer, K. E. Sex Differences in Lifespan, *Cell Metabolism*, **23**, no. 6 (2016), 1022–1033. DOI: [10.1016/j.cmet.2016.05.019](https://doi.org/10.1016/j.cmet.2016.05.019) (cit. on p. 4).
- [10] Barbara, M. A. et al. An Introduction to Epigenetics, *Neonatal Network*, **36**, no. 3 (May 2017), 124–128. DOI: [10.1891/0730-0832.36.3.124](https://doi.org/10.1891/0730-0832.36.3.124) (cit. on p. 6).
- [11] Barker, D. J. P. In utero programming of chronic disease, *Clinical Science*, **95**, no. 2 (Aug. 1998), 115–128. DOI: [10.1042/cs0950115](https://doi.org/10.1042/cs0950115) (cit. on p. 5).

- [12] Barker, D. J. P. The fetal and infant origins of adult disease. *BMJ: British Medical Journal*, **301**, no. 6761 (1990), 1111 (cit. on p. 5).
- [13] Barker, D. J. P. The origins of the developmental origins theory, *Journal of Internal Medicine*, **261**, no. 5 (May 2007), 412–417. DOI: [10.1111/j.1365-2796.2007.01809.x](https://doi.org/10.1111/j.1365-2796.2007.01809.x) (cit. on p. 6).
- [14] Barondes, S. *Molecules and Mental Illness*. Scientific American Library, 1999 (cit. on p. 4).
- [15] Barres, R. and Zierath, J. R. DNA methylation in metabolic disorders, *The American Journal of Clinical Nutrition*, **93**, no. 4 (Apr. 2011), 897S–900S. DOI: [10.3945/ajcn.110.001933](https://doi.org/10.3945/ajcn.110.001933) (cit. on p. 5).
- [16] Biganzoli, E. et al. Feed forward neural networks for the analysis of censored survival data: a partial logistic regression approach, *Statistics in Medicine*, **17**, no. 10 (1998), 1169–1186. DOI: [10.1002/\(SICI\)1097-0258\(19980530\)17:10<1169::AID-SIM796>3.0.CO;2-D](https://doi.org/10.1002/(SICI)1097-0258(19980530)17:10<1169::AID-SIM796>3.0.CO;2-D) (cit. on pp. 7, 25).
- [17] Blair, S. N. et al. Physical Fitness and All-Cause Mortality: A Prospective Study of Healthy Men and Women, *JAMA*, **262**, no. 17 (Nov. 1989), 2395–2401. DOI: [10.1001/jama.1989.03430170057028](https://doi.org/10.1001/jama.1989.03430170057028) (cit. on p. 4).
- [18] Boks, M. P. et al. The Relationship of DNA Methylation with Age, Gender and Genotype in Twins and Healthy Controls, *PLoS ONE*, **4**, no. 8 (Aug. 2009). Ed. by J. Najbauer, e6767. DOI: [10.1371/journal.pone.0006767](https://doi.org/10.1371/journal.pone.0006767) (cit. on p. 5).
- [19] Bookman, E. B. et al. Gene-environment interplay in common complex diseases: forging an integrative model—recommendations from an NIH workshop, *Genetic Epidemiology* (2011). DOI: [10.1002/gepi.20571](https://doi.org/10.1002/gepi.20571) (cit. on p. 4).
- [20] Bugliari, D. et al. *RAND HRS Data Documentation, Version P*. Tech. rep. 2016 (cit. on p. 16).
- [21] Commenges, D. Inference for multi-state models from interval-censored data, *Statistical Methods in Medical Research*, **11**, no. 2 (Apr. 2002), 167–182. DOI: [10.1191/0962280202sm279ra](https://doi.org/10.1191/0962280202sm279ra) (cit. on p. 15).
- [22] Conti, R. M. et al. “Early Retirement and DI/SSI Applications: Exploring the Impact of Depression”, *Health at Older Ages: The Causes and Consequences of Declining Disability among the Elderly*. NBER Chapters. National Bureau of Economic Research, Inc, July 2009, 381–408 (cit. on p. 6).
- [23] Crowther, M. J. and Lambert, P. C. Parametric multistate survival models: Flexible modelling allowing transition-specific distributions with application to estimating clinically useful measures of effect differences, *Statistics in Medicine*, **36**, no. 29 (2017), 4719–4742. DOI: [10.1002/sim.7448](https://doi.org/10.1002/sim.7448) (cit. on p. 7).

- [24] DiLoreto, R. and Murphy, C. T. The cell biology of aging, *Molecular Biology of the Cell*, **26**, no. 25 (Dec. 2015). Ed. by W. Bement, 4524–4531. DOI: [10.1091/mbc.e14-06-1084](https://doi.org/10.1091/mbc.e14-06-1084) (cit. on p. 5).
- [25] Elman, J. L. Finding structure in time, *Cognitive Science*, **14**, no. 2 (1990), 179–211. DOI: [10.1016/0364-0213\(90\)90002-E](https://doi.org/10.1016/0364-0213(90)90002-E) (cit. on p. 28).
- [26] Epel, E. S. et al. Accelerated telomere shortening in response to life stress, *Proceedings of the National Academy of Sciences*, **101**, no. 49 (2004), 17312–17315. DOI: [10.1073/pnas.0407162101](https://doi.org/10.1073/pnas.0407162101) (cit. on p. 5).
- [27] Esteller, M. Epigenetics in Cancer, *New England Journal of Medicine*, **358**, no. 11 (Mar. 2008), 1148–1159. DOI: [10.1056/nejmra072067](https://doi.org/10.1056/nejmra072067) (cit. on p. 5).
- [28] Faraggi, D. and Simon, R. A neural network model for survival data, *Statistics in Medicine*, **14**, no. 1 (Jan. 1995), 73–82. DOI: [10.1002/sim.4780140108](https://doi.org/10.1002/sim.4780140108) (cit. on pp. 7, 25).
- [29] Fernandez, T. et al. “Gaussian Processes for Survival Analysis”, *Advances in Neural Information Processing Systems 29*. Ed. by D. D. Lee et al. Curran Associates, Inc., 2016, 5021–5029 (cit. on p. 25).
- [30] Fisher, G. G. and Ryan, L. H. Overview of the Health and Retirement Study and Introduction to the Special Issue, *Work, Aging and Retirement*, **4**, no. 1 (Dec. 2017). Ed. by M. Wang, 1–9. DOI: [10.1093/workar/wax032](https://doi.org/10.1093/workar/wax032) (cit. on p. 16).
- [31] Fleming, T. R. Nonparametric Estimation for Nonhomogeneous Markov Processes in the Problem of Competing Risks, *The Annals of Statistics*, **6**, no. 5 (1978), 1057–1070 (cit. on pp. 7, 13, 14).
- [32] Fleming, T. R. and Harrington, D. P. *Counting Processes and Survival Analysis*. Wiley, 2005 (cit. on p. 15).
- [33] Gill, R. D. “Product-integration”, *Encyclopedia of Biostatistics*. American Cancer Society, 2005. DOI: [10.1002/0470011815.b2a11058](https://doi.org/10.1002/0470011815.b2a11058) (cit. on p. 13).
- [34] Gill, R. D. and Johansen, S. A Survey of Product-Integration with a View Toward Application in Survival Analysis, *The Annals of Statistics*, **18**, no. 4 (1990), 1501–1555 (cit. on p. 13).
- [35] Gluckman, P. D. et al. Effect of In Utero and Early-Life Conditions on Adult Health and Disease, *New England Journal of Medicine*, **359**, no. 1 (July 2008), 61–73. DOI: [10.1056/nejmra0708473](https://doi.org/10.1056/nejmra0708473) (cit. on p. 5).
- [36] Graves, A. *Supervised Sequence Labelling with Recurrent Neural Networks*. Springer Berlin Heidelberg, 2012. DOI: [10.1007/978-3-642-24797-2](https://doi.org/10.1007/978-3-642-24797-2) (cit. on pp. 8, 27).
- [37] Graves, A. et al. Neural Turing Machines, *CoRR*, **abs/1410.5401** (2014) (cit. on pp. 8, 27).

- [38] Hammer, B. On the approximation capability of recurrent neural networks, *Neurocomputing*, **31**, no. 1 (2000), 107–123. DOI: [https://doi.org/10.1016/S0925-2312\(99\)00174-5](https://doi.org/10.1016/S0925-2312(99)00174-5) (cit. on pp. 8, 27).
- [39] Hannum, G. et al. Genome-wide Methylation Profiles Reveal Quantitative Views of Human Aging Rates, *Molecular Cell*, **49**, no. 2 (2013), 359–367. DOI: [10.1016/j.molcel.2012.10.016](https://doi.org/10.1016/j.molcel.2012.10.016) (cit. on p. 5).
- [40] Harrell Frank E., J. et al. Evaluating the Yield of Medical Tests, *JAMA*, **247**, no. 18 (May 1982), 2543–2546. DOI: [10.1001/jama.1982.03320430047030](https://doi.org/10.1001/jama.1982.03320430047030) (cit. on p. 29).
- [41] Hayflick, L. The limited in vitro lifetime of human diploid cell strains, *Experimental Cell Research*, **37**, no. 3 (1965), 614–636. DOI: [10.1016/0014-4827\(65\)90211-9](https://doi.org/10.1016/0014-4827(65)90211-9) (cit. on p. 5).
- [42] Hochreiter, S. and Schmidhuber, J. Long Short-Term Memory, *Neural Computation*, **9**, no. 8 (1997), 1735–1780. DOI: [10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735) (cit. on pp. 8, 28).
- [43] Hornik, K. et al. Multilayer feedforward networks are universal approximators, *Neural Networks*, **2**, no. 5 (Jan. 1989), 359–366. DOI: [10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8) (cit. on pp. 7, 24).
- [44] Horvath, S. DNA methylation age of human tissues and cell types, *Genome Biology*, **14**, no. 10 (Dec. 2013), 3156. DOI: [10.1186/gb-2013-14-10-r115](https://doi.org/10.1186/gb-2013-14-10-r115) (cit. on p. 5).
- [45] Juster, F. T. and Suzman, R. An Overview of the Health and Retirement Study, *The Journal of Human Resources*, **30** (1995), S7. DOI: [10.2307/146277](https://doi.org/10.2307/146277) (cit. on p. 16).
- [46] Kanherkar, R. R. et al. Epigenetics across the human lifespan, *Frontiers in Cell and Developmental Biology*, **2**, no. 3 (Sept. 2014), 124–128. DOI: [10.3389/fcell.2014.00049](https://doi.org/10.3389/fcell.2014.00049) (cit. on p. 6).
- [47] Kaplan, E. L. and Meier, P. Nonparametric Estimation from Incomplete Observations, *Journal of the American Statistical Association*, **53**, no. 282 (1958), 457–481. DOI: [10.1080/01621459.1958.10501452](https://doi.org/10.1080/01621459.1958.10501452) (cit. on p. 13).
- [48] Karakus, M. C. and Patton, L. C. Depression and the Onset of Chronic Illness in Older Adults: A 12-Year Prospective Study, *The Journal of Behavioral Health Services & Research*, **38**, no. 3 (Feb. 2011), 373–382. DOI: [10.1007/s11414-011-9234-2](https://doi.org/10.1007/s11414-011-9234-2) (cit. on p. 6).
- [49] Katzman, J. L. et al. DeepSurv: personalized treatment recommender system using a Cox proportional hazards deep neural network, *BMC Medical Research Methodology*, **18**, no. 1 (Feb. 2018), 24. DOI: [10.1186/s12874-018-0482-1](https://doi.org/10.1186/s12874-018-0482-1) (cit. on pp. 7, 25).
- [50] Khoury, M. J. et al. Genome-Wide Association Studies, Field Synopses, and the Development of the Knowledge Base on Genetic Variation and Human Diseases, *American Journal of Epidemiology*, **170**, no. 3 (June 2009), 269–279. DOI: [10.1093/aje/kwp119](https://doi.org/10.1093/aje/kwp119) (cit. on p. 4).

- [51] Lee, C. et al. “Deephit: A deep learning approach to survival analysis with competing risks”, *Thirty-Second AAAI Conference on Artificial Intelligence*. 2018 (cit. on pp. 7, 22, 25, 31, 32).
- [52] Lipton, Z. C. et al. A critical review of recurrent neural networks for sequence learning, *arXiv preprint* (2015) (cit. on p. 27).
- [53] Ranganath, R. et al. “Deep Survival Analysis”, *Proceedings of the 1st Machine Learning for Healthcare Conference*. Vol. 56. PMLR. 2016, 101–114 (cit. on pp. 7, 25).
- [54] Raut, L. K. Exits from Disability: Estimates from a Competing Risk Model, *Social Security Bulletin*, 77, no. 3 (2017), 15–38 (cit. on p. 11).
- [55] Raut, L. K. Pathways to disability and mortality in mid-ages: Estimates from the Health and Retirement Study Data, *Draft Presented at the 2019 SGE Annual Conference* (2019) (cit. on p. 4).
- [56] Raut, L. K. Progression of health, mortality and morbidity in the aging process: genetics, epigenetics and fetal programming, (*mimeo*) *University of Chicago* (2019) (cit. on p. 4).
- [57] Ren, K. et al. “Deep recurrent survival analysis”, *Proc. AAAI*. 2019, 1–8 (cit. on pp. 7, 25).
- [58] Renna, F. Alcohol Abuse, Alcoholism, and Labor Market Outcomes: Looking for the Missing Link, *ILR Review*, 62, no. 1 (Oct. 2008), 92–103. DOI: [10.1177/001979390806200105](https://doi.org/10.1177/001979390806200105) (cit. on p. 6).
- [59] Seib, C. et al. Stress, Lifestyle, and Quality of Life in Midlife and Older Australian Women: Results From the Stress and the Health of Women Study, *Women’s Health Issues*, 24, no. 1 (2014), e43–e52. DOI: [10.1016/j.whi.2013.11.004](https://doi.org/10.1016/j.whi.2013.11.004) (cit. on p. 6).
- [60] Shalev, I. and Belsky, J. Early-life stress and reproductive cost: A two-hit developmental model of accelerated aging?, *Medical Hypotheses*, 90 (2016), 41–47. DOI: [10.1016/j.mehy.2016.03.002](https://doi.org/10.1016/j.mehy.2016.03.002) (cit. on p. 5).
- [61] Shalev, I., Entringer, S., et al. Stress and telomere biology: A lifespan perspective, *Psychoneuroendocrinology*, 38, no. 9 (2013), 1835–1842. DOI: [10.1016/j.psyneuen.2013.03.010](https://doi.org/10.1016/j.psyneuen.2013.03.010) (cit. on p. 5).
- [62] Siegelmann, H. T. and Sontag, E. D. “On the computational power of neural nets”, *Proceedings of the fifth annual workshop on Computational learning theory - COLT '92*. ACM Press, 1992. DOI: [10.1145/130385.130432](https://doi.org/10.1145/130385.130432) (cit. on pp. 8, 27).
- [63] Simons, R. L. et al. Economic hardship and biological weathering: The epigenetics of aging in a U.S. sample of black women, *Social Science & Medicine*, 150 (2016), 192–200. DOI: [10.1016/j.socscimed.2015.12.001](https://doi.org/10.1016/j.socscimed.2015.12.001) (cit. on pp. 5, 6).
- [64] Sonnega, A. et al. Cohort Profile: the Health and Retirement Study (HRS), *International Journal of Epidemiology*, 43, no. 2 (Mar. 2014), 576–585. DOI: [10.1093/ije/dyu067](https://doi.org/10.1093/ije/dyu067) (cit. on p. 16).

- [65] Steffick, D. E. Documentation of affective functioning measures in the Health and Retirement Study, *Ann Arbor, MI: University of Michigan* (2000) (cit. on p. 17).
- [66] Thornburg, K. L. et al. “In Utero Life and Epigenetic Predisposition for Disease”, *Epigenetics and Cancer, Part B*. Elsevier, 2010, 57–78. DOI: [10.1016/b978-0-12-380864-6.00003-1](https://doi.org/10.1016/b978-0-12-380864-6.00003-1) (cit. on p. 6).
- [67] Vaupel, J. W. Biodemography of human ageing, *Nature*, **464**, no. 7288 (Mar. 2010), 536–542. DOI: [10.1038/nature08984](https://doi.org/10.1038/nature08984) (cit. on p. 4).
- [68] Wreede, L. C. de et al. The mstate package for estimation and prediction in non- and semi-parametric multi-state and competing risks models, *Computer Methods and Programs in Biomedicine*, **99**, no. 3 (Sept. 2010), 261–274. DOI: [10.1016/j.cmpb.2010.01.001](https://doi.org/10.1016/j.cmpb.2010.01.001) (cit. on p. 15).
- [69] Zarulli, V. et al. Women live longer than men even during severe famines and epidemics, *Proceedings of the National Academy of Sciences*, **115**, no. 4 (2018), E832–E840. DOI: [10.1073/pnas.1701535115](https://doi.org/10.1073/pnas.1701535115) (cit. on p. 4).
- [70] Zhao, R. et al. Deep learning and its applications to machine health monitoring, *Mechanical Systems and Signal Processing*, **115** (2019), 213–237. DOI: [10.1016/j.ymssp.2018.05.050](https://doi.org/10.1016/j.ymssp.2018.05.050) (cit. on p. 25).