MATCH-Net: Dynamic Prediction in Survival Analysis using Convolutional Neural Networks

Daniel Jarrett*, Jinsung Yoon † , and Mihaela van der Schaar*, †

* Department of Engineering Science, University of Oxford, UK [†] Department of Electrical Engineering, University of California Los Angeles, USA daniel.jarrett@eng.ox.ac.uk, jsyoon0823@ucla.edu,mihaela@ee.ucla.edu

Abstract

Accurate prediction of disease trajectories is critical for early identification and timely treatment of patients at risk. Conventional methods in survival analysis are often constrained by strong parametric assumptions and limited in their ability to learn from high-dimensional data, while existing neural network models are not readily-adapted to the longitudinal setting. This paper develops a novel convolutional approach that addresses these drawbacks. We present MATCH-Net: a Missingness-Aware Temporal Convolutional Hitting-time Network, designed to capture temporal dependencies and heterogeneous interactions in covariate trajectories and patterns of missingness. To the best of our knowledge, this is the first investigation of temporal convolutions in the context of dynamic prediction for personalized risk prognosis. Using real-world data from the Alzheimer's Disease Neuroimaging Initiative, we demonstrate state-of-the-art performance without making any assumptions regarding underlying longitudinal or time-to-event processes—attesting to the model's potential utility in clinical decision support.

1 Introduction

In Alzheimer's disease—the annual cost of which exceeds \$800 billion globally [1]—the effectiveness of therapeutic treatments is often limited by the challenge of identifying patients at early enough stages of progression for treatments to be of use. As a result, accurate and personalized prognosis in earlier stages of cognitive decline is critical for effective intervention and subject selection in clinical trials. Conventional statistical methods in survival analysis often begin by choosing explicit functions to model the underlying stochastic process [2–7]. However, the constraints—such as linearity and proportionality [8, 9] in the popular Cox model—may not be valid or verifiable in practice.

In spite of active research, a conclusive understanding of Alzheimer's disease progression remains elusive, owing to heterogeneous biological pathways [10, 11], complex temporal patterns [12, 13], and diverse interactions [14, 15]. Hence Alzheimer's data is a prime venue for leveraging the potential advantages of deep learning models for survival. Neural networks offer versatile alternatives by virtue of their capacity as general-purpose function approximators, able to learn—without restrictive assumptions—the latent structure between an individual's prognostic factors and odds of survival.

Contributions. Our goal is to establish a novel convolutional model for survival prediction, using Alzheimer's disease as a case study for experimental validation. Primary contributions are threefold: First, we formulate a *generalized framework* for longitudinal survival prediction, laying the foundation for effective cross-model comparison. Second, our proposal is uniquely designed to capitalize on longitudinal data to issue *dynamically updated* survival predictions, accommodating potentially informative patterns of *data missingness*, and combining *human input* with model predictions. Third, we propose methods for deriving clinically meaningful insight into the model's inference process. Finally, we demonstrate state-of-the-art results in comparison with comprehensive benchmarks.

Machine Learning for Health (ML4H) Workshop at NeurIPS 2018.

2 Problem formulation

Let there be N patients in a study, indexed $i \in \{1, ..., N\}$. Time is treated as a discrete dimension of fixed resolution $\delta > 0$. Each longitudinal datum consists of the tuple $(t, \mathbf{x}_{i,t}, s_{i,t})$, where $\mathbf{x}_{i,t}$ is the vector of covariates recorded at time t, and $s_{i,t}$ is the binary survival indicator. Let random variable $T_{i,surv}$ denote the time-to-event, $T_{i,cens}$ the time of right-censoring, and $T_i = \min\{T_{i,surv}, T_{i,cens}\}$. Per convention, we assume that censoring is not correlated with the eventual outcome [16–18]. Let the complete longitudinal dataset be given by $\mathbf{X} = \{\langle (t, \mathbf{x}_{i,t}, s_{i,t}) \rangle_{i=0}^t \}_{i=1}^N$. Then we can define

$$\mathbf{X}_{i,t,w} = \langle (t', \mathbf{x}_{i,t'}, s_{i,t'}) \rangle_{t' \in T} \quad \text{where} \quad T = \{t' : t - w \le t' \le t\}$$
(1)

to be the set of observations for patient *i* extending from time *t* into a width-*w* window of the past, where parameter *w* depends on the model under consideration. Given longitudinal measurements in $\mathbf{X}_{i,t,w}$, our goal is to issue risk predictions corresponding to length- τ horizons into the future. Formally, given a *backward-looking* historical window (t - w, t], we are interested in the failure function for *forward-looking* prediction intervals $(t, t + \tau]$; that is, we want to estimate the probability

$$F_i(t+\tau|t,w) = \mathbb{P}(T_{i,\text{surv}} \le t+\tau|T_{i,\text{surv}} > t, \mathbf{X}_{i,t,w})$$
(2)

of event occurrence within each prediction interval. Observe that parameterizing the width of the historical window results in a generalized framework—for instance, a Cox model only utilizes the most recent measurement; that is, w = 1. At the other extreme, recurrent models may consume the entire history; that is, w = t. As we shall see, the best performance is in fact obtained via a flexible intermediate approach—that is, by learning the optimal width of a sliding window of history.



3 MATCH-Net

Figure 1: (a) The MATCH-Net architecture, with $\tau_{max} = 5\delta$. (b) Illustration of temporal convolutions acting over feature channels. (c) The longitudinal context within which MATCH-Net operates, as well as the network's prediction targets in association with the sliding window input mechanism.

We propose MATCH-Net: a Missingness-Aware Temporal Convolutional Hitting-time Network, innovating on current approaches in three respects. **Dynamic prediction**: Existing deep learning models issue prognoses on the basis of information from a single time point [19–29], potentially discarding valuable information in the presence of longitudinal data. We investigate *temporal convolutions* in capturing heterogeneous representations of temporal dependencies among observed covariates, enabling truly dynamic predictions on the basis of historical information. **Informative missingness**: Current survival methods rely on the common assumption that the timing and frequency of covariate measurements is uninformative [18, 30]. By contrast, our model is *missingness-aware*—that is, we explicitly account for informative missingness by learning correlations between patterns of missingness and disease progression. **Human input**: Instead of issuing predictions solely on the basis of patient state into model estimates to examine the incremental informativeness of *subjective* input. Each innovation is a source of gain in performance (see Section 4; further details in the appendix). MATCH-Net accepts as input a *sliding-window* of observed covariates $\mathbf{X}_{i,t,w}$, as well as a parallel binary mask of missing-value indicators $\mathbf{Z}_{i,t,w}$ taking on values of one to denote missing covariate measurements. In addition, the network optionally accepts a one-hot vector $\mathbf{r}_{i,t}$ describing the most recent clinician diagnosis of Alzheimer's disease progression. Starting from the base of the network, the convolutional block first learns representations of longitudinal covariate trajectories by extracting local features from temporal patterns. After each layer, filter activations from the auxiliary branch are concatenated with those in the main branch. Then, the fully-connected block captures more global relationships by combining local information. Finally, the output layers produce failure estimates

$$\hat{\mathbf{y}}_{i,t} = [F_i(t+\delta|t,w), ..., F_i(t+\tau_{\max}|t,w)]$$
(3)

for pre-specified prediction intervals, where τ_{max} is the maximal horizon desired. This *convolutional dual-stream* architecture explicitly captures representations of temporal dependencies within each stream, as well as between covariate trajectories and missingness patterns in association with disease progression. This accounts for the potential informativeness of both *irregular* sampling (*i.e.* intervals between consecutive clinical visits may vary) and *asynchronous* sampling (*i.e.* not all features are measured at the same time) [31, 32]. This also encourages the network to distinguish between actual measurements and imputed values, reducing sensitivity to specific imputation methods chosen. The final architecture uses convolutions of length $l = 5\delta$, with more filters per layer in the main branch.

Loss function. With the preceding notation, the negative log-likelihood of a single empirical result $s_{i,t+\tau}$ and model estimate $\hat{F}_i(t+\tau|t,w)$ in relation to some input window $\mathbf{X}_{i,t,w}$ is given by

$$l_{i,t,\tau}(\boldsymbol{\theta}) = -[s_{i,t+\tau} \log \hat{F}_i(t+\tau|t,w) + (1-s_{i,t+\tau})\log(1-\hat{F}_i(t+\tau|t,w))]$$
(4)

where θ denotes the parameters of the network. The total loss function is then computed to simultaneously take into account the quality of predictions for all prediction horizons τ , all times t available along each patient's longitudinal trajectory, and all patients $i \in \{1, ..., N\}$ in the survival dataset:

$$l(\boldsymbol{\theta}) = \frac{\delta^2}{\sum_{i=1}^N \sum_{j=1}^{t_i} \tau_i} \sum_{i=1}^N \sum_{j=1}^{t_i/\delta} \sum_{k=1}^{\tau_i/\delta} \alpha(i, (j-1)\delta, k\delta) \cdot l_{i, (j-1)\delta, k\delta}$$
(5)

where $\tau_i = \min\{t_i - t, \tau_{\max}\}$ accounts for failure or right-censoring. This is a natural generalization of the log-likelihood in [24] to accommodate longitudinal survival. Weight function $\alpha(i, t, \tau)$ allows trading off the relative importance of different patients, time steps, and prediction horizons. First, this allows standardizing patient contributions with $\alpha(i, t, \tau) \propto 1/t_i$, thereby counteracting the bias against patients with shorter survival durations. Second, in the context of heavily imbalanced classes, this allows up-weighting positive instances—that is, input windows that correspond to actual failure.

4 Experiments

The Alzheimer's Disease Neuroimaging Initiative (ADNI) study data is a longitudinal survival dataset of per-visit measurements for 1,737 patients [1]. The data tracks disease progression through clinical measurements at $\frac{1}{2}$ -year intervals, including quantitative biomarkers, cognitive tests, demographics, and risk factors. Our objective is to predict the *first stable occurrence* of Alzheimer's disease for each patient. Further information on the dataset, preparation, and training can be found in the appendix.

Benchmarks. We evaluate MATCH-Net against both traditional longitudinal methods in survival analysis and recent deep learning approaches; the former includes Cox landmarking and joint modeling methods, and the latter includes static and dynamic multilayer perceptrons and recurrent neural network models. Performance is evaluated on the basis of the area under the receiver operating characteristic curve (AUROC) and the area under the precision-recall curve (AUPRC), both computed with respect to prediction horizons τ . Five-fold cross validation metrics are reported in Table 1.

Performance. MATCH-Net produces state-of-the-art results, consistently outperforming both conventional statistical and neural network benchmarks. Gains are especially apparent in AUPRC scores—improving on the MLP by an average of 15% and on joint models by 16% across all horizons, and by 27% and 26% for one-step-ahead predictions. To understand the sources of improvement, we observe a 4% gain in AUPRC from introducing the sliding window mechanism (MLP to S-MLP), a 9% gain from incorporating temporal convolutions (S-MLP to S-TCN), and a further 2% gain from accommodating informative missingness (S-TCN to MATCH-Net). In addition, *including* the most recent clinician diagnosis results in a further 17% gain (further details located in the appendix).

Table 1: Cross validation performance for $\tau_{max} = 5\delta$ and $\delta = \frac{1}{2}$ years: MATCH-Net *without* clinician input, sliding-window temporal convolutional networks (S-TCN) and multilayer perceptrons (S-MLP). Benchmarks include fully-convolutional networks (FCN) [33] adapted for sequence-based survival prediction, Disease Atlas (D-Atlas) [34], baseline recurrent neural networks (RNN) including GRUs and LSTMs, static multilayer perceptrons (MLP), as well as conventional statistical methods for survival analysis—including joint modeling (JM) and Cox landmarking (LM). Bold values indicate best performance, and asterisks next to benchmark results indicate statistically significant difference (*p*-value < 0.05) from MATCH-Net result. More detailed breakdown of gains are found the appendix.

	τ	MATCH-Net	S-TCN	S-MLP	FCN	D-Atlas	RNN	MLP	JM	LM
AUROC	0.5	0.962	0.961	0.959	0.954	0.959	0.949*	0.948*	0.913*	0.909*
	1.0	0.942	0.941	0.932	0.930	0.929	0.930	0.930	0.917*	0.914*
	1.5	0.902	0.902	0.897	0.895	0.892	0.891	0.890	0.881	0.878
	2.0	0.909	0.908	0.904	0.903	0.896	0.901	0.895	0.894	0.890
	2.5	0.886	0.884	0.881	0.883	0.884	0.883	0.874	0.883	0.878
AUPRC	0.5	0.594	0.580	0.500	0.536	0.517	0.464*	0.469*	0.473*	0.469*
	1.0	0.513	0.505	0.447	0.453	0.423	0.410*	0.435	0.415*	0.412*
	1.5	0.373	0.367	0.354	0.357	0.364	0.340	0.340	0.319	0.325
	2.0	0.390	0.380	0.364	0.375	0.352	0.355	0.359	0.362	0.367
	2.5	0.384	0.381	0.371	0.365	0.360	0.365	0.356	0.366	0.363



Figure 2: Average saliency map indicating feature and temporal influence, computed using slopes of partial dependence on sample of numerical features across sliding window ($w = 5\delta$; $\delta = \frac{1}{2}$ years).

Visualization. From the preceding, we observe the largest gains by introducing convolutions. While this is consistent with our motivating hypothesis that convolutions are better able to capture temporal patterns, clinicians often desire a degree of transparency into the prediction process [35, 36]. We adopt the partial dependence approach in [37] to understand the input-output relationship, as well as examining the utility of convolutions. For each observed covariate *d*, we want to approximate how the estimated failure function varies based on the value of $\mathbf{x}_{t,w}^d$. We define the dependence

$$\mathbb{E}_{\mathbf{X}_{t,w}^{(-d)}}[\hat{F}(t+\tau|\mathbf{X}_{t,w})] \approx \frac{\delta}{\sum_{i=1}^{N} t_i} \sum_{i=1}^{N} \sum_{j=1}^{t_i/\delta} \hat{F}((j-1)\delta + \tau | \mathbf{x}_{(j-1)\delta,w}^d, \mathbf{X}_{(j-1)\delta,w}^{(-d)})$$
(6)

where $\mathbf{x}_{t,w}^d \cup \mathbf{X}_{t,w}^{(-d)} = \mathbf{X}_{t,w}$. By evaluating Equation 6 on the values $\mathbf{x}_{t,w}^d$ present in the data, the influence of each covariate can be measured by estimating its slope. For a global picture of what impact each feature and time step has on the model's predictions, we compute the influence for all features to produce an average *saliency map* [38, 39] highlighting the effect of convolutional layers. All else equal, absent convolutions we see that having worse covariate values any time step almost invariably has an upward impact on risk (*i.e.* negative impact on survival). On the other hand, with temporal convolutions we see that having worse covariate values at earlier time steps may result in a downward impact on risk (*i.e.* positive impact on survival), suggesting that convolutions may better facilitate modeling relative movements (*e.g.* sudden declines) than simply paying attention to levels. Further visualizations for added perspective on input-output relationships are found in the appendix.

References

- [1] Razvan V Marinescu, Neil P Oxtoby, Alexandra L Young, et al. Tadpole challenge: Prediction of longitudinal evolution in alzheimer's disease. *arXiv preprint arXiv:1805.03909*, 2018.
- [2] Kjell A Doksum and Arnljot Hbyland. Models for variable-stress accelerated life testing experiments based on wiener processes and the inverse gaussian distribution. *Technometrics*, 34(1):74–82, 1992.
- [3] D Kleinbaum and M Klein. Survival analysis statistics for biology and health. Survival, 510:91665, 2005.
- [4] Germán Rodríguez. Parametric survival models. Lectures Notes, Princeton University, 2005.
- [5] Mei-Ling Ting Lee and GA Whitmore. Proportional hazards and threshold regression: their theoretical and practical connections. *Lifetime data analysis*, 16(2):196–214, 2010.
- [6] Tamara Fernández, Nicolás Rivera, and Yee Whye Teh. Gaussian processes for survival analysis. In Advances in Neural Information Processing Systems, pages 5021–5029, 2016.
- [7] Ahmed M Alaa and Mihaela van der Schaar. Deep multi-task gaussian processes for survival analysis with competing risks. In Proceedings of the 30th Conference on Neural Information Processing Systems, 2017.
- [8] Ritesh Singh and Keshab Mukhopadhyay. Survival analysis in clinical trials: Basics and must know areas. *Perspectives in clinical research*, 2(4):145, 2011.
- [9] David R Cox. Regression models & life-tables. In Breakthroughs in statistics. Springer, 1992.
- [10] Laurel A Beckett, Michael C Donohue, Cathy Wang, et al. The alzheimer's disease neuroimaging initiative phase 2: Increasing the length, breadth, and depth of our understanding. *Alzheimer's & Dementia*, 11(7): 823–831, 2015.
- [11] John Hardy and Dennis J Selkoe. The amyloid hypothesis of alzheimer's disease: progress and problems on the road to therapeutics. *science*, 297(5580):353–356, 2002.
- [12] Bruno M Jedynak, Andrew Lang, Bo Liu, et al. A computational neurodegenerative disease progression score: method and results with the alzheimer's disease neuroimaging initiative cohort. *Neuroimage*, 63(3): 1478–1486, 2012.
- [13] Michael C Donohue, Hélène Jacqmin-Gadda, Mélanie Le Goff, et al. Estimating long-term multivariate progression from short-term data. *Alzheimer's & Dementia*, 10(5):S400–S410, 2014.
- [14] Lutz Frölich, Oliver Peters, Piotr Lewczuk, et al. Incremental value of biomarker combinations to predict progression of mild cognitive impairment to alzheimer's dementia. *Alzheimer's research & therapy*, 9(1): 84, 2017.
- [15] Tharick A Pascoal, Sulantha Mathotaarachchi, Monica Shin, et al. Synergistic interaction between amyloid and tau predicts the progression to dementia. *Alzheimer's & Dementia*, 13(6):644–653, 2017.
- [16] Anastasios A Tsiatis and Marie Davidian. Joint modeling of longitudinal and time-to-event data: an overview. *Statistica Sinica*, pages 809–834, 2004.
- [17] Yingye Zheng and Patrick J Heagerty. Partly conditional survival models for longitudinal data. *Biometrics*, 61(2):379–391, 2005.
- [18] Hans van Houwelingen and Hein Putter. Dynamic prediction in clinical survival analysis. CRC, 2011.
- [19] David Faraggi and Richard Simon. A neural network model for survival data. *Statistics in medicine*, 14(1): 73–82, 1995.
- [20] Jonathan Buckley and Ian James. Linear regression with censored data. *Biometrika*, 66(3):429–436, 1979.
- [21] Jared L Katzman, Uri Shaham, Alexander Cloninger, Jonathan Bates, Tingting Jiang, and Yuval Kluger. Deep survival: A deep cox proportional hazards network. *stat*, 1050:2, 2016.
- [22] Anny Xiang, Pablo Lapuerta, Alex Ryutov, Jonathan Buckley, and Stanley Azen. Comparison of the performance of neural network methods and cox regression for censored survival data. *Computational statistics & data analysis*, 34(2):243–257, 2000.
- [23] L Mariani, D Coradini, E Biganzoli, et al. Prognostic factors for metachronous contralateral breast cancer: a comparison of the linear cox regression model and its artificial neural network extension. *Breast cancer research and treatment*, 44(2):167–178, 1997.

- [24] Knut Liestbl, Per Kragh Andersen, and Ulrich Andersen. Survival analysis and neural nets. *Statistics in medicine*, 13(12):1189–1200, 1994.
- [25] Elia M Biganzoli, Federico Ambrogi, and Patrizia Boracchi. Partial logistic artificial neural networks (plann) for flexible modeling of censored survival data. In *Neural Networks, 2009. IJCNN 2009. International Joint Conference on*, pages 340–346. IEEE, 2009.
- [26] Chun-Nam Yu, Russell Greiner, Hsiu-Chin Lin, and Vickie Baracos. Learning patient-specific cancer survival distributions as a sequence of dependent regressors. In Advances in Neural Information Processing Systems, pages 1845–1853, 2011.
- [27] Margaux Luck, Tristan Sylvain, Héloïse Cardinal, Andrea Lodi, and Yoshua Bengio. Deep learning for patient-specific kidney graft survival analysis. arXiv preprint arXiv:1705.10245, 2017.
- [28] Stephane Fotso. Deep neural networks for survival analysis based on a multi-task framework. arXiv preprint arXiv:1801.05512, 2018.
- [29] Changhee Lee, William R Zame, Jinsung Yoon, and Mihaela van der Schaar. Deephit: A deep learning approach to survival analysis with competing risks. AAAI, 2018.
- [30] Dimitris Rizopoulos. Joint models for longitudinal and time-to-event data: With applications in R. Chapman and Hall/CRC, 2012.
- [31] Donald B Rubin. Inference and missing data. Biometrika, 63(3):581-592, 1976.
- [32] Seo-Jin Bang, Yuchuan Wang, and Yang Yang. Phased-lstm based predictive model for longitudinal ehr data with missing values.
- [33] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [34] Bryan Lim and Mihaela van der Schaar. Forecasting disease trajectories in alzheimer's disease using deep learning. arXiv preprint arXiv:1807.03159, 2018.
- [35] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144. ACM, 2016.
- [36] Anand Avati, Kenneth Jung, Stephanie Harman, Lance Downing, Andrew Ng, and Nigam H Shah. Improving palliative care with deep learning. In *Bioinformatics and Biomedicine (BIBM)*, 2017 IEEE International Conference on, pages 311–316. IEEE, 2017.
- [37] Jerome H Friedman. Greedy function approximation: a gradient boosting machine. Annals of statistics, pages 1189–1232, 2001.
- [38] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.
- [39] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. arXiv preprint arXiv:1312.6034, 2013.
- [40] Sarah Parisot, Sofia Ira Ktena, Enzo Ferrante, et al. Disease prediction using graph convolutional networks: Application to autism spectrum disorder and alzheimer's disease. *Medical image analysis*, 2018.
- [41] Kan Li and Sheng Luo. Functional joint model for longitudinal and time-to-event data: an application to alzheimer's disease. *Statistics in medicine*, 36(22):3560–3572, 2017.
- [42] Ke Liu, Kewei Chen, et al. Prediction of mild cognitive impairment conversion using a combination of independent component analysis and the cox model. *Frontiers in human neuroscience*, 11:33, 2017.
- [43] Stephanie JB Vos, Frans Verhey, Lutz Frölich, et al. Prevalence and prognosis of alzheimer's disease at the mild cognitive impairment stage. *Brain*, 138(5):1327–1338, 2015.
- [44] Terry M Therneau. A package for survival analysis in s. R package version 2.42, 2018.
- [45] FC Hu. Stepwise variable selection procedures for regression analysis. *R package version 0.1.0*, 2018.
- [46] Lang Wu. Mixed effects models for complex data. Chapman and Hall/CRC, 2009.
- [47] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [48] Takaya Saito and Marc Rehmsmeier. The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets. *PloS one*, 10(3):e0118432, 2015.

APPENDIX

Example use

While various settings may benefit from MATCH-Net as a matter of clinical decision support, we illustrate one possible application within the context of personalized screening. Figure 3 shows the historical risk trajectory and forward risk estimates for a randomly selected ADNI patient. During the first seven years of bi-annual visits, the patient exhibits cognitively normal behavior, and the MATCH-Net risk estimates—computed via measured biomarkers and neuropsychological tests—reflect this steady clinical state (see historical trajectory in blue). In fact, as of precisely seven years of follow-up, the predicted 30-month forward risk remains less than 10% (see predicted trajectory in orange).



Figure 3: Example application of MATCH-Net for personalized risk scoring.

However, two clinical visits later, when the patient returns for regular checkup and clinical measurements, the projected 30-month forward risk jumps to over 50% (see predicted trajectory in red). In this situation, the clinician is immediately alerted to the sudden increase in risk of dementia, and may decide to advise more frequent checkups, or to administer a wider range of tests and biomarker measurements in the immediate term to better assess the overall risk in light of the recent downturn. In fact, as it turns out in this case, the patient is indeed diagnosed with Alzheimer's disease at t = 10years, shedding light on MATCH-Net's potential as an early warning and subject selection system.

Related work

Model		Non- Linearity	Deep Learning	Direct-to- Probability	Time- Variance	Dynamic Prediction
Cox	(1972)	X	N/A	N/A	X	X
Faraggi & Simon	(1995)	1	X	X	X	X
Katzman et al.	(2016)	1	1	X	X	X
Luck et al.	(2017)	1	1	1	X	X
Lee et al.	(2018)	1	1	1	1	X
(This study)	(2018)	\checkmark	1	✓	\checkmark	\checkmark

Table 2: Summary of primary improvements by related work

The first study to investigate neural networks formally in the context of time-to-event analysis was done by [19]. By swapping out the linear functional in the Cox model for the topology of a hidden layer, their nonlinear proportional hazards approach was extended to other models for censored data, such as [4] and [20]. In 2016, [21] were the first to apply modern techniques in deep learning to survival, in particular without prior feature selection or domain expertise. While previous studies following [19]'s model generally produced mixed results [22, 23], [21] demonstrated comparable or superior performance of multilayer perceptrons in relation to conventional statistical methods.

Instead of predicting the hazard function as an intermediate objective, [24] first proposed—and [25] further developed—an alternative approach to predict survival directly for grouped time intervals. In 2017, [27] combined the use of the Cox partial likelihood with the goal of predicting probabilities for pre-specified time intervals. Inspired by the work of [26] on multi-task logistic regression models for survival, they generalized the idea to deep learning via multi-task multilayer perceptrons.

Recently, [29] and [28] in 2018 proposed learning the distribution of survival times directly, making no assumptions regarding the underlying stochastic processes—in particular with respect to the time-invariance of hazards. By being process-agnostic, [29] demonstrated significant improvements over existing statistical, machine learning, and neural network survival models on multiple real and synthetic datasets. In the context of Alzheimer's disease, [40] studied the use of medical image sequences for classifying patient progression. More pertinently, [34] used recurrent networks to forecast disease trajectories with ADNI data, but relied on the explicit assumption of exponential distributions. Finally, building on these developments, one of the main contributions of this study is the use of temporal convolutions for dynamic survival prediction—while making no assumptions, and allowing the associations between covariates and risks to evolve over time (see Table 2).

Details on dataset

We are primarily interested in the clinical status for each patient at any given time. An official diagnosis is recorded at each patient's visit, and consists of two attributes. First, each diagnosis may be either stable or transitive. The former consists of stable diagnoses of normal brain functioning ("NL"), mild cognitive impairment ("MCI"), or Alzheimer's disease ("AD"), and the latter consists of preliminary diagnoses indicating transitions between these categories, which may take the form of either conversions or reversions. Conversions indicate a forward progression in the disease trajectory, and reversions indicate a regression back towards an earlier stage of the disease.



Figure 4: State space of clinical diagnoses.

Patients may remain in stable or transition diagnosis states for any duration at a time. The average patient who receives a transition diagnosis is observed to persist in that state for one year, while some patients do not exit this state until almost 5 years have elapsed. Patients who receive a transition diagnosis may not actually be confirmed with a subsequent stable diagnosis; in fact, less than half of the transition diagnoses for dementia were confirmed by a stable diagnosis at the next step, and almost one quarter are never followed by a stable diagnosis at any point until right-censoring. In addition, patients often actually undergo reversion transitions back towards earlier stages of the disease; in fact, over 5% of the study population receive reversion diagnoses at some point in time.

Note on class imbalance. The per-patient failure rate is 14% (243 patients out of the total 1,737). However, given the online nature of the sliding window mechanism in training and testing, the effective fraction of observations with positive event labels for any prediction horizon is around 2%.

	Туре	Min	Max	Mean	S.D.	Missing
Event (AD)	Categorical	-	-	-	-	30.1%
Static						
Age	Numeric	5.4E+01	9.1E+01	7.4E+01	7.2E+00	0.0%
APOE4 (Risk)	Numeric	0.0E+00	2.0E+00	5.4E-01	6.6E-01	0.1%
Education Level	Numeric	4.0E+00	2.0E+01	1.6E+01	2.9E+00	0.0%
Ethnicity	Categorical	-	-	-	-	0.0%
Gender	Categorical	-	-	-	-	0.0%
Marital Status	Categorical	-	-	-	-	0.0%
Race	Categorical	-	-	-	-	0.0%
Biomarker						
Entorhinal	Numeric	1.0E+03	6.7E+03	3.4E+03	8.1E+02	49.2%
Fusiform	Numeric	7.7E+03	3.0E+04	1.7E+04	2.8E+03	49.2%
Hippocampus	Numeric	2.2E+03	1.1E+04	6.7E+03	1.2E+03	46.6%
Intracranial	Numeric	2.9E+02	2.1E+06	1.5E+06	1.7E+05	37.6%
Mid Temp	Numeric	8.0E+03	3.2E+04	1.9E+04	3.1E+03	49.2%
Ventricles	Numeric	5.7E+03	1.6E+05	4.2E+04	2.3E+04	41.6%
Whole Brain	Numeric	6.5E+05	1.5E+06	1.0E+06	1.1E+05	39.7%
Cognitive						
ADAS (11-item)	Numeric	0.0E+00	7.0E+01	1.1E+01	8.6E+00	30.1%
ADAS (13-item)	Numeric	0.0E+00	8.5E+01	1.8E+01	1.2E+01	30.7%
CRD Sum of Boxes	Numeric	0.0E+00	1.8E+01	2.2E+00	2.8E+00	29.7%
Mini Mental State	Numeric	0.0E+00	3.0E+01	2.7E+01	4.0E+00	29.9%
RAVLT Forgetting	Numeric	-1.2E+01	1.5E+01	4.2E+00	2.5E+00	30.9%
RAVLT Immediate	Numeric	0.0E+00	7.5E+01	3.5E+01	1.4E+01	30.7%
RAVLT Learning	Numeric	-5.0E+00	1.4E+01	4.0E+00	2.8E+00	30.7%
RAVLT Percent	Numeric	-5.0E+02	1.0E+02	6.0E+01	3.8E+01	31.4%

Table 3: Summary and description of variables used in ADNI dataset.

Data preparation

Since the ADNI dataset is an amalgamation from multiple related studies, most features are sparsely populated. Features with less than half of the entries missing are retained, leaving 18 numeric and 4 categorical features (see Table 3); the latter are represented by one-hot encoding, resulting in 16 binary features. Consistent with existing Alzheimer's studies, patients are aligned according to time elapsed since baseline measurements [41–43]. Timestamps are mapped onto a discrete axis with a fixed resolution of $\delta = \frac{1}{2}$ -year intervals; where multiple measurements qualify for the same destination, the most recent measurement per feature takes precedence. Original measurements were made at roughly $\frac{1}{2}$ -year intervals, so we observe that the average absolute deviation between original values and final timestamps amounts to an insignificant 4 days (*i.e.* less than 2% of each interval).

Where measurements are missing, values are reconstructed using zero-order hold interpolation. In addition, due to the fixed-width nature of the sliding window, the input tensor $\mathbf{X}_{i,t,w}$ for initial prediction times $t < w - \delta$ correspond to left-truncated information $t - w + \delta < 0$; feature values are therefore extrapolated backwards for all intervals of the form $[-w + \delta, -\delta]$. Note that regardless of the imputation mechanism, information on original patterns of missingness—due to truncation, irregular sampling, and asynchronous sampling alike [31, 32]—is preserved in the missing-value mask $\mathbf{Z}_{i,t,w}$ provided in parallel to the network. Finally, to improve numerical conditioning, all features are normalized with empirical means and standard deviations from the training set data.

Training procedure

Training begins with input tuple $(\mathbf{X}, \mathbf{Z}, \mathbf{R})$, where $\mathbf{X} = \{\langle \mathbf{X}_{i,t,w} \rangle_{t=0}^{t_i}\}_{i=1}^N$, $\mathbf{Z} = \{\langle \mathbf{Z}_{i,t,w} \rangle_{t=0}^{t_i}\}_{i=1}^N$, and $\mathbf{R} = \{\langle \mathbf{r}_{i,t} \rangle_{t=0}^{t_i}\}_{i=1}^N$, and terminates with a set of calibrated network weights $\boldsymbol{\theta}$. The network is trained until convergence, up to a maximum of 50 epochs. Training loss is only computed for

event labels corresponding to actual recorded clinical visits (*i.e.* timestamps with recorded covariate values); neither imputed nor forward-filled labels are included. Analogous to our total loss function definition, the convergence metric is the sum of performance scores across all prediction tasks,

$$C = \sum_{k=1}^{\tau_{\max}/\delta} \beta(k\delta) \cdot \text{AUROC}_{k\delta} + \gamma(k\delta) \cdot \text{AUPRC}_{k\delta}$$
(7)

where $\beta(\tau)$, $\gamma(\tau)$ optionally allow trading off the relative importance between the two measures and different horizons. In this study we simply use the unweighted sum, although any convex combination would be valid. Empirically, results are not meaningfully improved by favoring one metric over the other. Elastic net regularization is used, and validation performance is computed every 10 iterations. For early stopping, validation scores serve as proxies for generalization error. Positive instances are oversampled to counteract class imbalance. To augment training data, artificial labels for dummy events of the form $\tau = 0$ are generated, and positive labels are filled forward for all horizons $\tau > \tau_0$ where τ_0 corresponds to the first failure for that example. Patients diagnosed with Alzheimer's at baseline (20%) are excluded from testing and validation, since survival is undefined in those cases.

Table 4: Hyperparameter selection ranges for random search.

Hyperparameter	Selection Range
Connected Layers	1, 2, 3, 4, 5
Convolutional Layers	1, 2, 3, 4, 5
Dropout Rate	0.1, 0.2, 0.3, 0.4, 0.5
Epochs for Convergence	1, 2, 3, 4, 5, 6, 7, 8, 9, 10
Learning Rate	1e-4, 3e-4, 1e-3, 3e-3, 1e-2, 3e-2
L1-Regularlisation	None, 1e-4, 3e-4, 1e-3, 3e-3, 1e-2, 3e-2, 1e-1
L2-Regularization	None, 1e-4, 3e-4, 1e-3, 3e-3, 1e-2, 3e-2, 1e-1
Minibatch Size	32, 64, 128, 256, 512
Number of Filters (Covariates)	32, 64, 128, 256, 512
Number of Filters (Masks)	8, 16, 32, 64, 128
Oversample Ratio	None, 1, 2, 3, 5, 10
Recurrent Unit State Size	$1 \times, 2 \times, 3 \times, 4 \times, 5 \times$
Width of Connected Layers	32, 64, 128, 256, 512
Width of Convolutional Filters	3, 4, 5, 6, 7, 8, 9, 10
Width of Sliding Window	3, 4, 5, 6, 7, 8, 9, 10

For Cox landmarking, we use the implementation in [44] for interval-censored data, fitting a sequence of proportional hazards regression models for observation groups. Optimal groupings are determined by exhaustive search in $\frac{1}{2}$ -year increments. Preliminary feature-selection is performed by stepwise regression for the best candidate model, using the implementation of [45]. For joint modeling, we adopt the common two-stage method in [46], first fitting linear mixed effects sub-models for significant variables, then fitting landmarking models based on mean estimates from the sub-models. In both cases, consistent with literature, time is defined as years since initial follow-up [41–43].

For all neural network models, hyperparameter optimization is carried out via 100 iterations of random search(see Table 4). In addition, activation functions are also searched over for gating units in recurrent network cells. Model selection is performed on the basis of final composite scores—as in Equation 7—for each candidate. We use 5-fold cross validation to evaluate performance, stratified at the patient level—that is, patients are randomly selected into datasets for training (60%), validation (20%), and testing (20%), with the ratio of positive patients (*i.e.* those for which at least one sliding window contains at least one outcome of failure) to negative patients is kept uniform across folds.

Further analysis

In addition, we utilize *output scatters* to aid our understanding of the input-output relationship. We use MC dropout to generate average responses by explicitly varying one input feature at a time [47] in equation 6. This is done by either (1) varying only the final value of the measurement within the sliding window, or by (2) varying all values of the measurement the same time. While the latter gives

insight into how the response changes with respect to a difference in *levels*, the former sheds light on how the response is affected by a sudden *change* instead. Interestingly in the case of temporal convolutions, for multiple features (Figure 5 shows the MMSE and CDRSB features as examples) the response is actually *stronger* to final value changes. This effect is not present in any feature for multilayer perceptrons. In light of the performance of convolutional models, this is consistent with our hypothesis that they appear able to learn more complex temporal patterns than multilayer perceptrons; in this case, these results suggest that a patient who experiences a sudden decline in certain features is expected to fare worse than one who has always scored just as poorly all along.



Figure 5: Comparison of output scatters for changes in all values within sliding window (red) versus changes in the final value only (blue). Convolutional models show stronger responses in the latter.

Sources of gain

While the advantage of using multilayer perceptrons over traditional statistical survival models has been studied (see [29] for example), here we account for the additional sources of gain from our design choices. Table 5 (a) shows the initial benefit from incorporating *longitudinal histories* of covariate measurements. Recurrent networks are used as a reasonable starting point; improvements-where positive—are marginal at best. (b) However, experiments on differenced data—capturing relative movements without absolute levels—indicate substantial informativeness (see df-RNN trained on differences vs. joint models trained on levels). (c) Utilizing limited sliding windows instead of recurrent cells shows promising improvements, boosting average AUPRC by 4% over both RNN and MLP models. (d) The addition of temporal convolutions produces incremental AUPRC gains of 9%. (e) Accommodating informative missingness by introducing the dual-stream architecture results in further incremental AUPRC gains of 2%. Compared with the joint modeling baseline, MATCH-Net achieves average AUPRC improvements of 15% and one-step-ahead improvements of 26%. (f) Furthermore, incorporating the most recent clinician input (MATCH-Net⁺) improves AUPRC by an additional 17%. While the distribution of gains skews near-term with our default choices of α , β , and γ (*i.e.* unweighted by τ), alternate distributions may be obtained via appropriate weight functions depending on deployment context. (g) Finally, using MATCH-Net as an example, we observe individual and cumulative benefits due to miscellaneous design choices applied to all models.

Table 5: Source-of-ga	ain accounting. Bo	ld values indicate	best performance.	Note that AUROC is
much less sensitive th	an AUPRC in the co	ontext of highly im	balanced classes lik	e the ADNI data [48].

	(a) (Gain from	Covariate I	History		(b) Info	ormativenes	s of Diffe	rences			
	$\mid \tau$	R	NN	Ν	1LP	 df-	RNN	J	M			
AUROC	0.5	0.949	±0.009	0.948	±0.010	0.913	±0.026	0.913	±0.036			
	1.0	0.930	± 0.012	0.930	± 0.011	0.830	± 0.020	0.917	± 0.016			
	1.5	0.891	± 0.026	0.890	± 0.027	0.700	± 0.057	0.881	± 0.029			
	2.0	0.901	± 0.025	0.895	± 0.029	0.712	± 0.059	0.894	± 0.034			
	2.5	0.883	± 0.031	0.874	± 0.039	0.708	± 0.083	0.883	± 0.036			
AUPRC	0.5	0.464	± 0.079	0.469	±0.064	0.494	±0.057	0.473	± 0.072			
	1.0	0.410	± 0.060	0.435	± 0.056	0.386	± 0.036	0.415	± 0.042			
	1.5	0.340	± 0.067	0.340	± 0.067	0.213	± 0.058	0.319	± 0.044			
	2.0	0.355	± 0.068	0.359	± 0.065	0.241	± 0.078	0.362	± 0.050			
	2.5	0.365	± 0.087	0.356	± 0.085	0.217	± 0.096	0.366	± 0.064			
	(c)	Gain fron	n Limited W	⁷ indow		 (d) Gain	from Temp	poral Convolutions S-MLP				
	$\mid \tau$	S-]	MLP	R	.NN	 S-'	ГCN	S-]	MLP			
AUROC	0.5	0.950	±0.009	0.949	± 0.009	0.961	±0.005	0.950	± 0.009			
	1.0	0.932	±0.012	0.930	± 0.012	0.941	± 0.007	0.932	± 0.012			
	1.5	0.897	± 0.025	0.891	± 0.026	0.902	± 0.025	0.897	± 0.025			
	2.0	0.904	± 0.026	0.901	± 0.025	0.908	± 0.026	0.904	± 0.026			
	2.5	0.881	± 0.035	0.883	± 0.031	0.884	± 0.032	0.881	± 0.035			
AUPRC	0.5	0.500	±0.066	0.464	±0.079	0.580	±0.066	0.500	±0.066			
	1.0	0.447	± 0.056	0.410	± 0.060	0.505	± 0.065	0.447	± 0.056			
	1.5	0.354	± 0.061	0.340	± 0.067	0.367	± 0.063	0.354	± 0.061			
	2.0	0.364	± 0.054	0.355	± 0.068	0.380	± 0.052	0.364	± 0.054			
	2.5	0.371	± 0.084	0.365	± 0.087	0.381	± 0.085	0.371	± 0.084			
	(e) Gai	n from M	issingness-A	Awareness	5	 (f) Gain	from Most	st Recent Diagnosis				
	$\mid \tau$	Mat	CH-Net	S-'	TCN	MATC	MATCH-Net ⁺ MATCH-N					
AUROC	0.5	0.962	±0.004	0.961	± 0.005	0.989	±0.003	0.962	± 0.004			
	1.0	0.942	± 0.007	0.941	± 0.007	0.951	± 0.008	0.942	± 0.007			
	1.5	0.902	± 0.024	0.902	± 0.025	0.897	± 0.026	0.902	± 0.024			
	2.0	0.909	± 0.027	0.908	± 0.026	0.901	± 0.024	0.909	± 0.027			
	2.5	0.886	± 0.033	0.884	± 0.032	0.885	± 0.024	0.886	± 0.033			
AUPRC	0.5	0.594	±0.058	0.580	± 0.066	0.862	± 0.037	0.594	± 0.058			
	1.0	0.513	± 0.059	0.505	± 0.065	0.636	± 0.043	0.513	± 0.059			
	1.5	0.373	± 0.065	0.367	± 0.063	0.393	± 0.082	0.373	± 0.065			
	2.0	0.390	± 0.059	0.380	± 0.052	0.380	± 0.051	0.390	± 0.059			
	2.5	0.384	± 0.081	0.381	± 0.085	0.375	± 0.066	0.384	± 0.081			
		-				 						

(S) Guin from oversumpting (G), faber for warding (E), and elastic fiet (E) for infinence	(g)	Gain f	rom oversamp	oling (O), label	forwarding	(L), and	l elastic n	et (E) fo	r Match-N
---	-----	--------	--------------	----------	----------	------------	----------	-------------	-----------	-----------

	AUROC						AUPRC				
ΟLΕ	0.5	1.0	1.5	2.0	2.5	0.5	1.0	1.5	2.0	2.5	
ХХХ	0.946	0.923	0.895	0.902	0.888	0.493	0.418	0.348	0.351	0.355	
ХХХ	0.956	0.938	0.901	0.906	0.886	0.534	0.477	0.368	0.381	0.384	
XXX	0.944	0.920	0.888	0.896	0.886	0.435	0.413	0.334	0.351	0.357	
XJJ	0.961	0.941	0.902	0.908	0.887	0.575	0.502	0.370	0.384	0.377	
$\checkmark \times \times$	0.946	0.923	0.894	0.902	0.887	0.499	0.425	0.352	0.365	0.365	
\checkmark \checkmark \checkmark	0.956	0.938	0.901	0.907	0.887	0.533	0.473	0.362	0.379	0.385	
\checkmark \checkmark \times	0.942	0.918	0.885	0.894	0.881	0.448	0.404	0.332	0.348	0.361	
\checkmark	0.962	0.942	0.902	0.909	0.886	0.594	0.513	0.373	0.390	0.384	