

Forecasting Disease Trajectories in Alzheimer’s Disease Using Deep Learning

Bryan Lim
University of Oxford
bryan.lim@eng.ox.ac.uk

Mihaela van der Schaar
University of Oxford
Alan Turing Institute
mihaela.vanderschaar@eng.ox.ac.uk

ABSTRACT

Joint models for longitudinal and time-to-event data are commonly used in longitudinal studies to forecast disease trajectories over time. Despite the many advantages of joint modeling, the standard forms suffer from limitations that arise from a fixed model specification and computational difficulties when applied to large datasets. We adopt a deep learning approach to address these limitations, enhancing existing methods with the flexibility and scalability of deep neural networks while retaining the benefits of joint modeling. Using data from the Alzheimer’s Disease Neuroimaging Institute¹, we show improvements in performance and scalability compared to traditional methods.

ACM Reference Format:

Bryan Lim and Mihaela van der Schaar. 2018. Forecasting Disease Trajectories in Alzheimer’s Disease Using Deep Learning. In *Proceedings of KDD Workshop on Machine Learning for Medicine and Healthcare (KDD ’18)*. London, UK, 4 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 OVERVIEW

Effective clinical decision support often involves the dynamic forecasting of medical conditions based on clinically relevant variables collected over time. This involves jointly predicting the expected time to events of interest (e.g. death), biomarker trajectories, and other associated risks at different stages of disease progression.

With the prevalence of aging populations around the globe, Alzheimer’s disease (AD) is a significant threat to public health - growing from being relatively rare at the start of the 20th century to having a case being reported every 7 seconds around the world [1]. Patients at risk of developing Alzheimer’s disease are usually monitored over time based on longitudinal cognitive scores and MRI measurements [6], which help doctors evaluate the severity of a patient’s condition and formulate a diagnosis. As such, the ability to produce joint forecasts - such as those depicted in Figure 1 -

would help doctors determine both the likelihood of developing dementia and the expected rate of deterioration of a given patient, potentially allowing for intervention at an early stage.

Given the promising initial results found in our companion paper for Cystic Fibrosis patients [5], we investigate the application of the Disease-Atlas - a novel conception of the joint modeling framework using deep learning - to jointly predicting the expected time-to-transition to Alzheimer’s disease and the values of longitudinal measurements, providing additional clinical decision support to doctors evaluating potential Alzheimer’s disease patients. We start with an overview of the Disease-Atlas in Section 2 and 3, demonstrating performance gains for tests on data from the ADNI in Section 4.

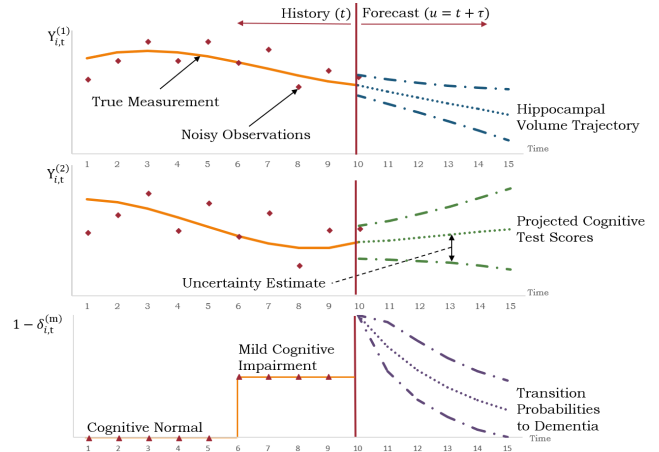


Figure 1: Disease-Atlas Predictions over Time

2 PROBLEM DEFINITION

For a given longitudinal study, let there be N patients with observations made at time t , for $0 \leq t \leq T_{cens}$ where T_{cens} denotes an administrative censoring time². For the i^{th} patient at time t , observations are made for a K -dimensional vector of longitudinal variables $V_{i,t} = [Y_{i,t}^{(1)}, \dots, Y_{i,t}^{(C)}, B_{i,t}^{(1)}, \dots, B_{i,t}^{(D)}]$, where $Y_{i,t}^{(c)}$ and $B_{i,t}^{(d)}$ are continuous and discrete longitudinal measurements respectively, a L -dimensional vector of external covariates $X_{i,t} = [X_{i,t}^{(1)}, \dots, X_{i,t}^{(L)}]$, and a M -dimensional vector of event occurrences $\delta_{i,t} = [\delta_{i,t}^{(1)}, \dots, \delta_{i,t}^{(M)}]$, where $\delta_{i,t}^{(m)} \in \{0, 1\}$ is an indicator

² Administrative censoring refers to the right-censoring that occurs when a study observation period ends.

¹For the Alzheimer’s Disease Neuroimaging Initiative: Data used in preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

KDD ’18, August 2018, London, UK

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

variable denoting the presence or absence of the m^{th} event. $T_{i,t}^{(m)}$ is defined to be the first time the event is observed after t , which allows us to model both repeated events and events that lead to censoring (e.g. death). The final observation for patient i occurs at $T_{i,\max} = \min(T_{cens}, T_{i,0}^{(a_1)}, \dots, T_{i,0}^{(a_{\max})})$, where $\{a_1, \dots, a_{\max}\}$ is the set of indices for events that censor observations.

3 DISEASE-ATLAS ARCHITECTURE

The Disease-Atlas captures the associations within the joint modeling framework, by learning *shared representations* between trajectories at different stages of the network, while retaining the same sub-model distributions captured by joint models. The network, as shown in Figure 2, is conceptually divided into 3 sections: 1) A shared temporal layer to learn the temporal and cross-correlations between variables, 2) task-specific layers to learn shared representations between related trajectories, and 3) an output layer which computes parameters for predictive sub-model distributions for use in likelihood loss computations during training and generating predictive distributions at run-time.

The equations for each layer are listed in detail below. For notational convenience, we drop the subscript i for variables in this section, noting that the network is only applied to trajectories from one patient at time. While we focus on both the continuous valued and time-to-event predictions for tests in Section 4, we include descriptions of binary predictions here for completeness.

Shared Temporal Layer. We start with an RNN at the base of the network, which incorporates historical information into forecasts by updating its memory state over time. For the tests in Section 4, we adopt the use of a long-short term memory network (LSTM) in the base layer.

$$[\mathbf{h}_t, \mathbf{m}_t] = \text{RNN}([\mathbf{X}_t, \mathbf{V}_t], \mathbf{m}_{t-1}) \quad (1)$$

Where \mathbf{h}_t is the output of the RNN and \mathbf{m}_t its memory state. To generate uncertainty estimates for forecasts and retain consistency with joint models, we adopt the MC dropout approach described in [3]. Dropout masks are applied to the inputs, memory states and outputs of the RNN, and are also fixed across time steps. For memory updates, the RNN uses the Exponential Linear Unit (ELU) activation function.

Task-specific Layers. For the task-specific layers, variables can be grouped according to the types of outputs, with layer $\mathbf{z}_{c,t}$ for continuous-valued longitudinal variables, $\mathbf{z}_{b,t}$ for binary longitudinal variables and $\mathbf{z}_{e,t}$ for events. Dropout masks are also applied to the outputs of each layer here. At the inputs to the continuous and binary task layers, a prediction horizon τ is also concatenated with the outputs from the RNN. This allows the parameters of the predictive distributions at $t + \tau$ to be computed in the final layer, i.e. $\tilde{\mathbf{h}}_t = [\mathbf{h}_t, \tau]$.

$$\mathbf{z}_{c,t} = \text{ELU}(\mathbf{W}_c \tilde{\mathbf{h}}_t + \mathbf{a}_c) \quad (2a)$$

$$\mathbf{z}_{b,t} = \text{ELU}(\mathbf{W}_b \tilde{\mathbf{h}}_t + \mathbf{a}_b) \quad (2b)$$

$$\mathbf{z}_{e,t} = \text{ELU}(\mathbf{W}_e \tilde{\mathbf{h}}_t + \mathbf{a}_e) \quad (2c)$$

Output Layer. The final layer computes the parameter vectors of the predictive distribution, which are used to compute log likelihoods during training and dynamic predictions at run-time.

$$\mu_{t+\tau} = \mathbf{W}_\mu \mathbf{z}_{c,t} + \mathbf{a}_\mu \quad (3a)$$

$$\sigma_{t+\tau} = \text{Softplus}(\mathbf{W}_\sigma \mathbf{z}_{c,t} + \mathbf{a}_\sigma) \quad (3b)$$

$$\mathbf{p}_{t+\tau} = \text{Sigmoid}(\mathbf{W}_p \mathbf{z}_{b,t} + \mathbf{a}_p) \quad (3c)$$

$$\lambda_t = \text{Softplus}(\mathbf{W}_\lambda \mathbf{z}_{e,t} + \mathbf{a}_\lambda) \quad (3d)$$

Softplus activation functions are applied to $\sigma_{t+\tau}$ and $\mathbf{p}_{t+\tau}$ to ensure that we obtain valid (i.e. ≥ 0) standard deviations and binary probabilities. For simplicity, the exponential distribution is selected to model survival times, and predictive distributions are as below:

$$Y_{t+\tau}^{(c)} \sim N\left(\mu_{t+\tau}^{(c)}, \sigma_{t+\tau}^{(c)2}\right) \quad (4a)$$

$$B_{t+\tau}^{(d)} \sim \text{Bernoulli}\left(p_{t+\tau}^{(d)}\right) \quad (4b)$$

$$T_t^{(m)} \sim \text{Exponential}\left(\lambda_t^{(m)}\right) \quad (4c)$$

3.1 Multitask Learning

From the above, the negative log-likelihood of the data given the network is:

$$\begin{aligned} \mathcal{L}(\mathbf{W}) = \sum_{i,t,w,k_c,k_b,m} & - \left[\log f_c\left(Y_{i,t+\tau}^{(c)} | \mu_{t+\tau}^{(c)}, \sigma_{t+\tau}^{(c)2}, \mathbf{W}\right) \right. \\ & + \log f_b\left(B_{i,t+\tau}^{(d)} | p_{t+\tau}^{(d)}, \mathbf{W}\right) \\ & \left. + \log f_T\left(T_{i,t}^{(m)} | \lambda_t^{(m)}, \mathbf{W}\right) \right] \quad (5) \end{aligned} \quad (6)$$

Where $f_c(\cdot), f_b(\cdot)$ are likelihood functions based on Equations 4 and \mathbf{W} collectively represents the weights and biases of the entire network. For survival times, $f_T(\cdot)$ is given as:

$$f_T\left(T_t^{(m)} | \lambda_t^{(m)}, \mathbf{W}\right) = \left(\lambda_t^{(m)}\right)^{\delta_{i,T}} \exp\left(-\lambda_t^{(m)} T_t^{(m)}\right) \quad (7)$$

Which corresponds to event-free survival until time T before encountering the event [2]. While the negative log-likelihood can be directly optimized across tasks, the use of multitask learning can yield the following benefits:

Better Survival Representations. As shown in [4], multitask learning problems which have one main task of interest can weight the individual loss contributions of each subtask to favor representations for the main problem. For our current architecture, where we group similar tasks into task-specific layers, our loss function corresponds to:

$$\begin{aligned} L(\mathbf{W}) = & - \underbrace{\alpha_c \sum_{i,t,w,c} \log f_c\left(Y_{t+\tau}^{(c)} | \mathbf{W}\right)}_{\text{Continuous Longitudinal Loss } l_c} - \underbrace{\alpha_b \sum_{i,t,w,d} \log f_b\left(B_{t+\tau}^{(d)} | \mathbf{W}\right)}_{\text{Binary Longitudinal Loss } l_b} \\ & - \underbrace{\alpha_T \sum_{i,t,m} \log f_T\left(T_t^{(m)} | \mathbf{W}\right)}_{\text{Time-to-event Loss } l_T} \quad (8) \end{aligned}$$

Given that survival predictions are the primary focus of many longitudinal studies, we set $\alpha_c = \alpha_b = 1$ and include α_T as an additional hyperparameter to be optimized. To train the network,

Algorithm 1 Training Disease-Atlas

Input: Data $\Omega = \{\Omega_1, \dots, \Omega_Q\}$, max iterations \mathcal{J}
Output: Calibrated network weights \mathbf{W}
for count= 1 **to** \mathcal{J} **do**
 Get minibatch $\mathcal{M} \sim \gamma$ random samples from Ω
 Sample task loss function $l \sim \{l_c, l_b, l_T\}$
 Update $\mathbf{W} \leftarrow \text{Adam}(l, \mathcal{M})$, using feed-forward passes with dropout applied
end for

patient trajectories are subdivided into Q sets of $\Omega_q(i, \rho, \tau) = \{\mathbf{X}_{i,0:\rho}, \mathbf{Y}_{i,\rho+\tau}, \mathbf{T}_{\max,i}, \delta_i\}$, where ρ is the length of the covariate history to use in training trajectories up to a maximum of ρ_{\max} . Full details on the procedure can be found in Algorithm 1.

Handling Irregularly Sampled Data. We address issues with irregular sampling by grouping variables that are measured together into the same task, and training the network with multitask learning. For instance, volumes of different parts of the brain (e.g. hippocampal, ventricular and intra-cranial volume) that are measured together during the same MRI scan session can be grouped together in the same task. Given the completeness of the datasets we consider, we assume that task groupings match those defined by the task-specific layer of the network, and multitask learning is performed using Equation 8 and Algorithm 1.

We note, however, that in the extreme case where none of the trajectories are aligned, we can define each variable as a separate task with its own loss function l_* . Algorithm 1 then samples loss functions for one variable at a time, and the network is trained using only actual observations as target labels. This could reduce errors in cases where multiple sample rates exist and simple imputation is used, which might result in the multioutput networks replicating the imputation process instead of making true predictions.

3.2 Forecasting Disease Trajectories

Dynamic prediction involves 2 key elements - 1) calculating the expected longitudinal values and survival curves as described above, and 2) computing uncertainty estimates. To obtain these measures, we apply the Monte-Carlo dropout approach of [3] by approximating the posterior over network weights as:

$$p(V_{t+\tau}^{(k)} | \mathcal{F}_t) \approx \frac{1}{J} \sum_{j=1}^J p(V_{t+\tau}^{(k)} | \mathcal{F}_t, \hat{\mathbf{W}}_j) \quad (9)$$

Where we draw J samples $\hat{\mathbf{W}}_j$ using feed-forward passes through the network with the same dropout mask applied across time-steps. The samples obtained can then be used to compute expectations and uncertainty intervals for forecasts.

4 PERFORMANCE EVALUATION FOR ALZHEIMER’S DISEASE

4.1 Data Description

The Alzheimer’s Disease Neuroimaging Initiative (ADNI) study data is a comprehensive dataset that tracks the progression of the Alzheimer’s disease (AD) through 3 main states: normal brain function, mild cognitive impairment and the onset of either the disease

or dementia. This data surveys 1737 patients for periods up to 10 years, capturing informative features extracted with Positron Emission Tomography (PET) regions of interest (ROI) scans – e.g. measures of cell metabolism, which are known to be reduced for AD patients – Magnetic Resonance (MRI) and Diffusion Tensor imaging (DTI) (for instance, ventricles volume), CSF and blood biomarkers, genetics, cognitive tests (ADAS-Cog), demographic and others. Observations were discretized to 6-month (or 0.5 year) intervals, and missing measurements were imputed using the previous value if present, and the population mean otherwise. In this investigation, we use a random selection of 60% of patients for our training data, 20% for validation and the final 20% for evaluation as per the CF tests. This was repeated 3 times to form 3 different partitions of the dataset, which were then used for cross-validation. The Disease-Atlas was used to jointly forecast longitudinal observations of clinical scores and scan measurements, treating the transition to Alzheimer’s Disease from either mild cognitive impairment (MCI) or Cognitively Normal (CN) states as our event of interest. Hyperparameter optimization was performed with 20 iterations of random search.

4.2 Results & Discussion

To evaluate predictions of the event-of-interest – i.e. transitions to dementia – we compared the performance of the Disease-Atlas against simpler recurrent neural networks (i.e. LSTMs) and standard methods from biostatistics (i.e. landmarking [7] and joint models (JM) fitted with a two-step approximation [8]).

Prediction results for transitions to dementia – in terms of the area under the receiver operating characteristic (AUROC) and the precision-recall curve (AUPRC) – and MSE improvements for longitudinal forecasts can be found in Tables 1 and 2 respectively. From the cross-validation performance, we see that the Disease-Atlas consistently outperforms both the standard neural network and traditional benchmarks for survival analysis particularly on short-term horizons – improving the LSTM by 10% and JM by 7% on average across all time steps.

For longitudinal predictions, we focus on both the Disease-Atlas and JM which are able to generate predictions at arbitrary time steps in the future. Once again, the Disease-Atlas outperforms joint models across the majority of longitudinal variables and time steps, with gains of 40% on average – highlighting the benefits of a deep learning approach to joint modeling.

5 CONCLUSIONS

In this paper, we investigate an application of the Disease-Atlas to forecasting longitudinal measurements and expected time-to-transition to Dementia for patients at risk of Alzheimer’s Disease. Using data from the ADNI, the Disease-Atlas [5] demonstrated performance gains over both simpler neural networks such as LSTMs and traditional methods from biostatistics – demonstrating the advantages of the Disease-Atlas as a method for joint modeling and highlighting its potential as a tool for clinical decision support.




Figure 2: Disease-Atlas Network Architecture

Table 1: Cross-Validation Performance for Transitions to Alzheimer’s Disease (Mean \pm S.D.)

	τ	Disease-Atlas	LSTM
AUROC	0.5	0.954 (\pm 0.008)	0.938 (\pm 0.005)
	1	0.935 (\pm 0.006)	0.929 (\pm 0.005)
	1.5	0.906 (\pm 0.004)	0.905 (\pm 0.001)
	2	0.899 (\pm 0.014)	0.899 (\pm 0.006)
	τ	Landmarking	JM
	0.5	0.913 (\pm 0.033)	0.916 (\pm 0.035)
	1	0.914 (\pm 0.010)	0.919 (\pm 0.016)
	1.5	0.892 (\pm 0.012)	0.897 (\pm 0.007)
	2	0.884 (\pm 0.023)	0.890 (\pm 0.015)
	τ	Disease-Atlas	LSTM
AUPRC	0.5	0.326 (\pm 0.038)	0.256 (\pm 0.036)
	1	0.271 (\pm 0.043)	0.268 (\pm 0.015)
	1.5	0.211 (\pm 0.043)	0.198 (\pm 0.018)
	2	0.183 (\pm 0.056)	0.178 (\pm 0.027)
	τ	Landmarking	JM
	0.5	0.270 (\pm 0.048)	0.295 (\pm 0.078)
	1	0.240 (\pm 0.056)	0.257 (\pm 0.083)
	1.5	0.174 (\pm 0.040)	0.185 (\pm 0.040)
	2	0.167 (\pm 0.050)	0.183 (\pm 0.045)

REFERENCES

- [1] Gavril Cornutiu. 2015. The Epidemiological Scale of Alzheimer’s Disease. *Journal of Clinical Medicine Research* 7, 9 (2015), 657–666.
- [2] George H. Duntzman and Moon-Ho R. Ho. 2006. *An Introduction to Generalized Linear Models*. SAGE Publications, Inc., Thousand Oaks, CA, USA.
- [3] Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in Neural Information Processing Systems (NIPS 2016)*.
- [4] F Li, L Tran, K-H Thung, S Ji, D Shen, and J. Li. 2015. A Robust Deep Model for Improved Classification of AD/MCI Patients. *IEEE journal of biomedical and health informatics*. 19, 5 (2015), 1610–1616.

Table 2: % Decrease in MSE for Longitudinal Predictions between Disease-Atlas & JM (Mean \pm S.D.)

τ (Years)	0.5	1	1.5	2
MRI				
ICV	63% (\pm 8%)	62% (\pm 8%)	62% (\pm 7%)	61% (\pm 6%)
WholeBrain	51% (\pm 9%)	51% (\pm 10%)	51% (\pm 10%)	51% (\pm 8%)
Ventricles	78% (\pm 7%)	77% (\pm 7%)	77% (\pm 6%)	76% (\pm 6%)
Hippocampus	63% (\pm 0%)	63% (\pm 0%)	62% (\pm 1%)	62% (\pm 1%)
Fusiform	61% (\pm 12%)	59% (\pm 12%)	57% (\pm 12%)	56% (\pm 11%)
MidTemp	64% (\pm 5%)	62% (\pm 5%)	61% (\pm 4%)	59% (\pm 4%)
Entorhinal	56% (\pm 9%)	54% (\pm 7%)	52% (\pm 6%)	50% (\pm 6%)
Cognitive				
CDRSB	34% (\pm 4%)	30% (\pm 2%)	27% (\pm 2%)	25% (\pm 2%)
MMSE	18% (\pm 3%)	16% (\pm 1%)	16% (\pm 1%)	17% (\pm 2%)
ADAS11	-8% (\pm 6%)	9% (\pm 5%)	15% (\pm 6%)	19% (\pm 5%)
RAVLT Imm.	39% (\pm 4%)	36% (\pm 2%)	33% (\pm 2%)	31% (\pm 2%)
RAVLT Learn.	21% (\pm 1%)	20% (\pm 1%)	19% (\pm 2%)	18% (\pm 3%)
ADAS13	0% (\pm 11%)	14% (\pm 10%)	19% (\pm 9%)	22% (\pm 10%)
RAVLT Forget.	23% (\pm 4%)	32% (\pm 5%)	32% (\pm 5%)	33% (\pm 6%)
RAVLT Perc. F.	13% (\pm 5%)	29% (\pm 4%)	32% (\pm 3%)	35% (\pm 0%)

- [5] B. Lim and M. van der Schaar. 2018. Disease-Atlas: Navigating Disease Trajectories with Deep Learning. In *Proceedings of the 3rd Machine Learning for Healthcare Conference (MLHC 2018)*.
- [6] Judith Neugroschl and Sophia Wang. 2011. Alzheimer’s Disease: Diagnosis and Treatment Across the Spectrum of Disease Severity. *The Mount Sinai Journal of Medicine* 78, 4 (2011), 596–612.
- [7] Hans van Houwelingen and Hein Putter. 2011. *Dynamic Prediction in Clinical Survival Analysis*. CRC Press, Inc., Boca Raton, FL, USA.
- [8] Lang Wu. 2009. *Mixed Effects Models for Complex Data*. Chapman & Hall/CRC, Hoboken, NJ, USA.