

Title: EHR-based Patient Risk Stratification Tool for Probable AD

Donna Tjandra¹, Raymond Migrino², Bruno Giordani³, and Jenna Wiens¹

¹Division of Computer Science and Engineering, University of Michigan, Ann Arbor, MI

²Phoenix Veterans Affairs Health Care System, Phoenix, AZ

³Neuropsychology Section, Department of Psychiatry, University of Michigan, Ann Arbor, MI

Background:

To date, predictive modeling for Alzheimer's disease (AD) risk has focused on data not routinely collected in clinical care and is limited to short prediction horizons (e.g., 2-4 years). Given the limitations of existing datasets, we sought to leverage electronic health records (EHRs) that contain decades of longitudinal data for thousands of patients, with the goal of developing and validating a predictive model for AD onset with a 10-year prediction horizon.

Methods:

A retrospective study included patients aged 68-72 admitted to Michigan Medicine prior to 2008 with at least 10 years of follow-up. To control for age, we aligned patients between 68-72 years. Given EHR data up to and including the visit of alignment, we sought to predict the primary outcome: probable AD within the next 10 years (labeled by a previously validated cohort discovery tool). Patients were repeatedly randomly divided into a training and test sets. Using EHR-based covariates (e.g., demographics, diagnoses, procedures, vital signs, laboratory results, medications), we learned a linear model to predict AD. Predictive performance was measured on held-out test data using area under the receiver operating characteristics curve (AUROC). Covariates were ranked according to predictive power, using permutation importance.

Results:

After applying inclusion/exclusion criteria (Figure 1), patients in the final study population (n=8,416) had a median of 11 encounters prior to alignment (interquartile range (IQR)=4-25) and a median of 84 encounters during follow-up (IQR=36-172). The model achieved an AUROC of 0.708 (IQR=0.686-0.733). Important predictive factors included healthcare utilization, testing for secondary causes of AD, and cardiovascular factors (Figure 2). Identifying at-risk patients more than 6 years in advance was possible, though more difficult (Figure 3).

Conclusions:

Using longitudinal EHR data, we can predict AD in advance of clinical diagnosis with modest accuracy. Mining routinely collected data could shed light on AD progression, especially in the decades before clinical onset.

Figure 1: Applying Inclusion/Exclusion Criteria. We begin with all patients in Michigan Medicine’s Research Data Warehouse (RDW). Numbers in each box correspond to the number of patients included/excluded. To simplify our analysis, we exclude censored patients and require that patients have at least 10 years of follow-up post alignment, unless they convert to probable AD sooner.

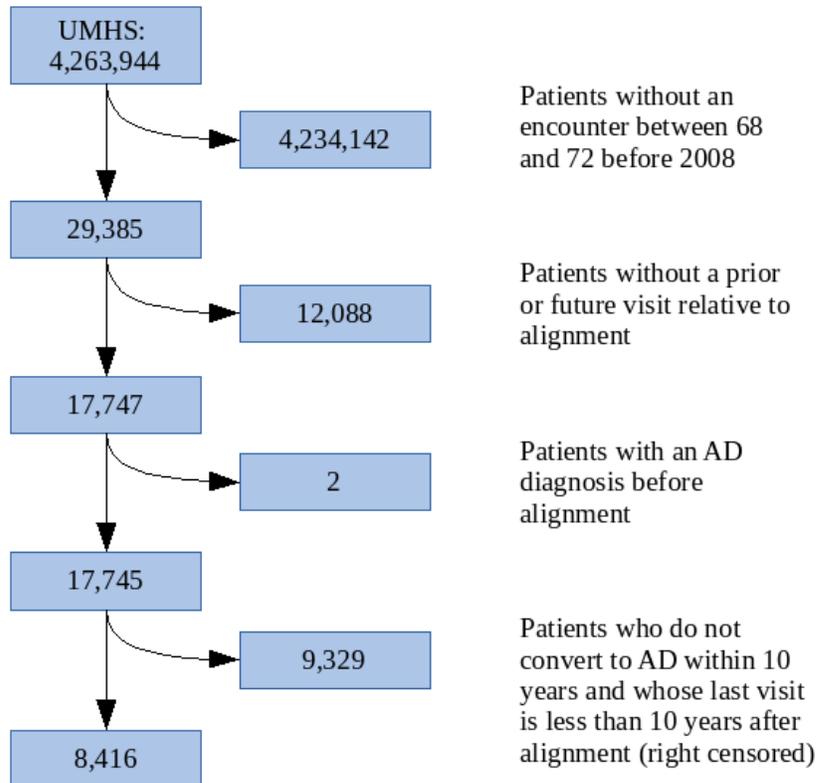


Figure 2: Top 10 most predictive features. We analyzed the predictive power of our features using permutation importance to identify which ones affected the AUROC the most. Results reported are averaged over 100 permutations, and error bars represent 95% confidence intervals. Risk factors correspond to features whose weights were positive, and protective factors (marked with an *) correspond to features whose weights were negative.

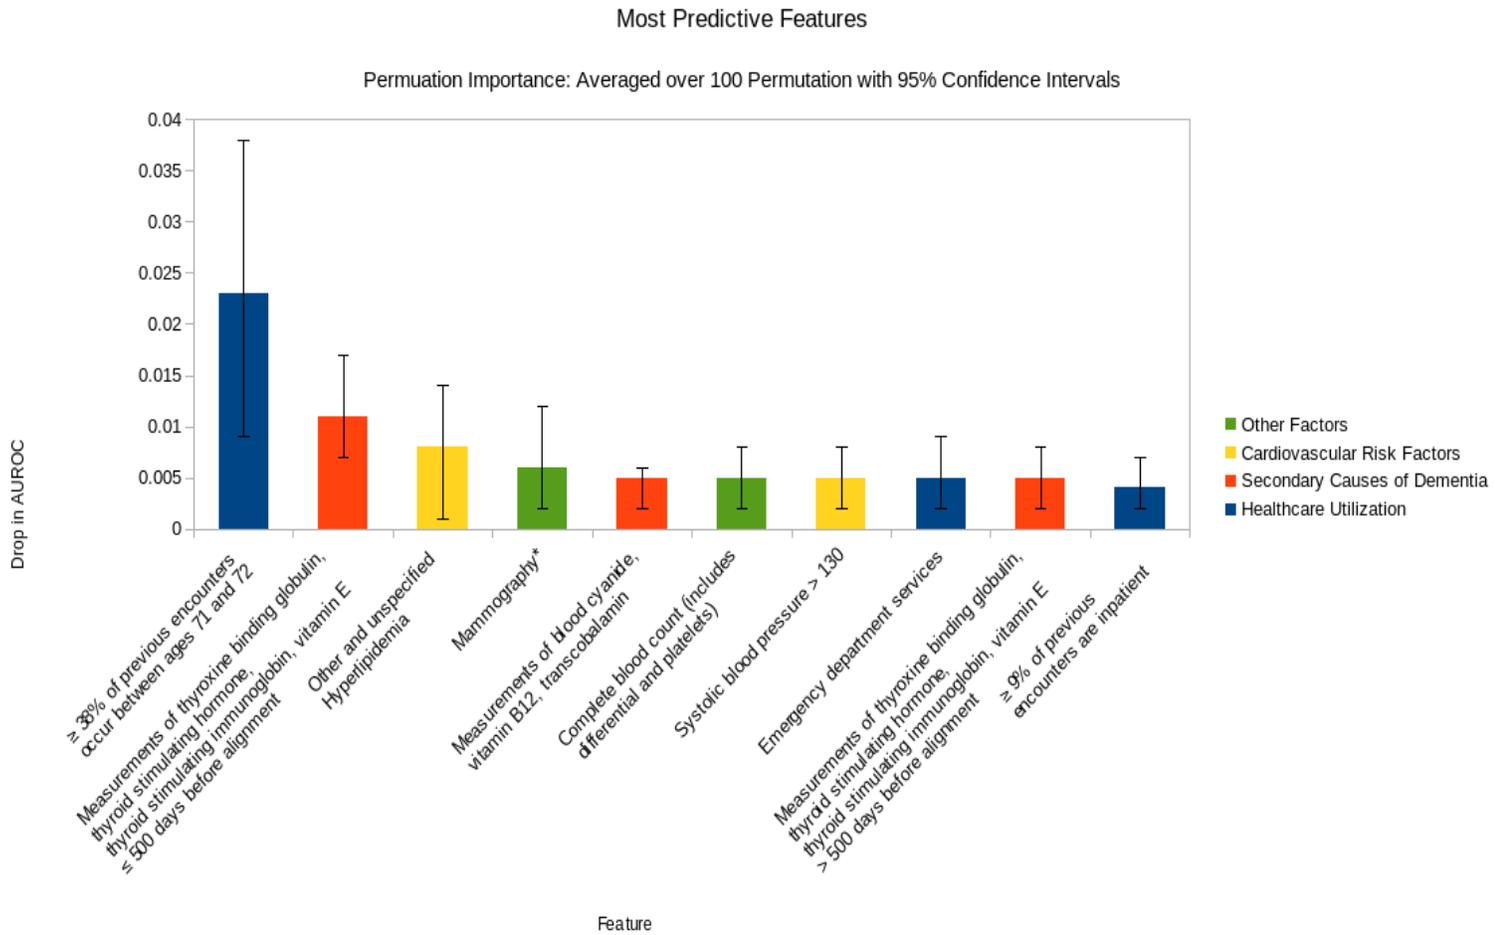


Figure 3: Model validation. Among those identified to convert to probable AD, we analyzed the model's ability to predict conversion relative to amount of time it took to convert starting from alignment. Patients were predicted to convert if the probability of conversion given by the model was above the 65th percentile. Our model was able to predict conversion on large time windows as well as small ones. Error bars represent 95% confidence intervals over 1,000 bootstrapped samples.

