Resting State Functional MRI Based Alzheimer's Disease Classification

Final Project

ECE 784 Computer Vision

Michael Evans

Due: 30 April 2026

Submitted: 30 April 2026

Abstract.

Resting state functional MRI (rs-fMRI) measures spontaneous BOLD signal fluctuations across the brain and is a useful modality for studying neurodegenerative disease. In this report, we use rs-fMRI to classify Alzheimer's Disease (AD) from cognitively normal (CN) subjects on a subset of the ADNI cohort. We first preprocess each subject with fMRIPrep and parcellate the brain into 120 regions of interest with the AAL2 atlas. We then compute four resting state metrics per region of interest: mean BOLD intensity, functional connectivity (FC) with Fisher z transform, fractional amplitude of low-frequency fluctuations (fALFF), and regional homogeneity (ReHo). Next, we train a class-balanced linear support vector machine (SVM) to classify AD versus CN. On 49 AD and 17 CN subjects with 7,620 features, the linear SVM which uses all features achieves a class balanced accuracy of 0.710 under stratified 5-fold cross-validation and 0.649 under leave-one-out cross-validation, with F1 (AD) of 0.787 and 0.699 respectively. We also use the SelectKBest algorithm for dimensionality reduction and feature importance to find that the K = 200 features reaches the highest balanced accuracy of 0.682, indicating that heavy feature reduction can still generate comparable results with using the full feature vector on this small dataset.

Technical Discussion.

Resting state functional MRI captures slow, spontaneous fluctuations in the BOLD signal while the subject is at rest in the scanner [1]. To extract usable features from each subject, we first preprocess every BOLD run with fMRIPrep [2] in MNI152NLin2009cAsym space, filter out volumes with framewise displacement above 0.5 mm [3] for motion correction, and bandpass the cleaned signal to 0.01–0.08 Hz. fMRIPrep performs intensity non-uniformity correction, skull stripping, and registration. Volumes filtered out for motion correction are cubic spline interpolated from surrounding volumes to make up this removal. Voxels are then aggregated into 120 regions of interest (ROIs) using the AAL2 atlas [4]. Two filtering passes are run per subject: a regression only pass that preserves the full spectrum (used for ALFF and fALFF) and a regression with bandpass filter pass (used for FC, mean BOLD, and ReHo). We then compute the four resting state metrics as defined below.

Mean BOLD. For each voxel v with preprocessed time series $S_v(t)$ of length T, the mean BOLD intensity $\overline{S_v}(t)$ is the temporal mean signal:

$$\overline{S_v}(t) = \frac{1}{T}\sum_{i=1}^{T} S_v(t) \qquad 1$$

Voxel-wise means are then averaged within each ROI R using the AAL2 label image to give one ROI-level feature per subject:

$$\mu_R = \frac{1}{|R|}\sum\nolimits_{v\in R} \mu_v \qquad 2$$

Functional Connectivity. Functional connectivity (FC) quantifies the temporal coupling of signal fluctuations between two ROI time series $x_i(t)$ and $x_j(t)$ after bandpass filtering. To measure the correlation between each ROI we use the Pearson correlation coefficient:

$$r_{ij} = \frac{\sum_{t=1}^{n}(x_i(t) - \overline{x_i}) \times (x_j(t) - \overline{x_j})}{\sqrt{\sum_{t=1}^{n}(x_i(t) - \overline{x_i})^2 \times \sum_{t=1}^{n}\left(x_j(t) - \overline{x_j}\right)^2}} \qquad 3$$

The Pearson coefficient r is bounded on [−1, 1], and to stabilize the variance across subjects we apply the Fisher z transform [5] to create a normally distributed form of the Pearson FC matrix:

$$z' = \frac{1}{2}\ln\left(\frac{1+r}{1-r}\right) \quad (4)$$

Which results in a 120×120 z-transformed FC matrix. This Fisher z matrix is then flattened along its upper triangle to yield 7,140 ROI-pair connectivity features per subject.

ALFF and fALFF. The amplitude of low-frequency fluctuations (ALFF) and its fractional form (fALFF) measure how much energy a voxel carries within a narrow, neurologically meaningful 0.01-0.08 Hz band [6]. For each voxel we compute the discrete Fourier transform of the regression-only (non-bandpassed) signal:

$$X_v(f) = \sum_{t=1}^{T} s_v(t) \times e^{-j \times 2\pi f t/T} \quad (5)$$

ALFF is then just the sum of FFT amplitudes in the specified low-frequency band 0.01-0.08 Hz:

$$ALFF_v = \frac{1}{N}\sum_{i=1}^{N}\sqrt{S(f_i)} \quad (6)$$

Where $S(f_i)$ is the power at frequency f_i. fALFF then normalizes ALFF by the total spectral amplitude across the full spectrum, which reduces the effect of noise [7]:

$$fALFF_v = \frac{ALFF_v}{\sum_{\forall f}|X_v(f)|} \quad (7)$$

Regional Homogeneity (ReHo). ReHo measures how synchronized a voxel's time series is with its immediate neighbors using Kendall's coefficient of concordance W [8]. For a $3 \times 3 \times 3$ cluster of $k = 27$ voxels around a target voxel, we can rank the bandpassed signal at each timepoint, sum the ranks across the neighborhood N, $R_t = \sum_{v \in N} rank(s_v(t))$, and compare against the expected mean, $\overline{R} = \frac{k(T+1)}{2}$ which is defined by W:

$$W = 12 \times \frac{\sum_{t=1}^{T}\left(R_t - \overline{R}\right)^2}{k^2(T^3 - T)} \quad (8)$$

W is also bounded on $[0,1]$, where 1 indicates perfect synchronization of the signals across the region. We compute W voxel-wise and then average within each AAL2 ROI to yield a 120-

dimensional ReHo feature vector per subject, representing how the similarly the voxels in each region fluctuate with each other through time.

Classifier. For each subject we concatenate the 120 mean BOLD, ALFF, fALFF, ReHo, and 7,140 upper-triangular Fisher-z FC features into a single feature vector of length 7,620. We then train a simple class-balanced linear support vector machine (SVM) [9] using scikit-learn [10]. The SVM architecture was chosen for this task due its known ability to work well on high dimensional feature spaces and low number of samples. We also train an SVM on a reduced set of selected features using the SelectKBest selection algorithm [10]. Models are evaluated under both stratified 5-fold and leave one out cross validation on $K \in \{25, 50, 100, 200, 500, 1000\}$.

Discussion of Results.

We use 71 ADNI [11] subjects with usable preprocessed resting-state BOLD runs from fMRIPrep, split as 53 AD and 18 CN. After dropping rows with any NaN feature, 66 subjects remained (49 AD, 17 CN), giving a ~3:1 class imbalance that we reduce the effects of at training time with SVM class weighting. The group averaged figures in [A1-A7] are computed from an averaging of the features obtained from equations 1-8 on 53 AD and 18 CN subjects.

Figures [A1] and [A2] show the mean BOLD per ROI for the AD and CN groups. The general shape across ROIs is the same in both groups, but AD shows higher absolute intensities in most regions, which was unexpected. This most likely reflects scanner scaling and site differences in the ADNI cohort rather than a clinical finding. These effects could be mitigated with scanner harmonization, a method not included in this report. Mean BOLD on its own is therefore likely not a strong discriminator between AD and CN.

Figures [A3] and [A4] show the per ROI fALFF for both groups. The overall shape is again similar, but AD is visibly shifted lower than CN across the entire spectrum, with several regions showing a notable decrease around the ROIs 40, 60, and 120. This is consistent with the expected reduction in low-frequency BOLD power in AD and suggests that fALFF is a much more informative feature than mean BOLD signal, which is expected.

In figures [A5] and [A6] we show the Fisher z FC matrix averaged within each group. The global pattern is comparable, with a strong block of within-network connectivity around ROIs 45-60. Figure [A7] is the AD-CN difference matrix on a more sensitive color scale: AD shows broadly

stronger long-range connectivity than CN (predominantly red), but there is a clear cluster of negative values (blue) in the same ROI 45-60 block where CN exceeds AD. This suggests the useful disease signal in FC is concentrated in a specific network rather than spread evenly across the brain, and that local ROI analysis is probably less informative than network-level analysis for this task. While the higher average global connectivity in the AD FC matrix is unexpected, the clustering to a network-level difference is expected.

Figures [B1-B3] summarize the classifier performance across feature selection settings under both 5-fold and leave-one-out cross-validation. The all-features linear SVM and the SelectKBest K = 200 configuration are the two strongest performing classifiers across all three metrics. The all-features model achieves a balanced accuracy of 0.710 under 5-fold and 0.649 under leave-one-out, with an F1 score (AD) of 0.787 and 0.699 and F1 (CN) of 0.558 and 0.490 respectively. SelectKBest with K = 200 under leave-one-out notably achieves the highest balanced accuracy of 0.682, with F1 (AD) of 0.809 and F1 (CN) of 0.526. Intermediate K values (25, 100, 500, 1000) consistently underperform, suggesting that either the full multimodal feature vector or a strong ANOVA filter both outperform a weaker feature filter.

Across all configurations, the F1 score for AD as the positive class is substantially higher than for the CN group, which is a consequence of the large class imbalance. To mitigate this effect, we use a class-balanced loss in the model to penalize misclassifications of the minority class more heavily. In conclusion, we extracted four resting-state functional MRI metrics per AAL2 ROI, plus a Fisher-z FC matrix from a small ADNI cohort preprocessed with fMRIPrep to classify AD versus CN with a linear SVM. The best configurations reach a balanced accuracy in the high 0.6 to low 0.7 range. With the larger class, AD, we achieve an F1 score of close to 0.8, and an F1 score of close to 0.6 with CN, indicating that performance would greatly benefit from more data. The most notable future steps for this work are to increase size and reduce class imbalance. Moreover, summarizing features at the network level rather than the ROI level are likely to increase performance, and correspond stronger to the established literature.

[This space intentionally left blank.]

Results.

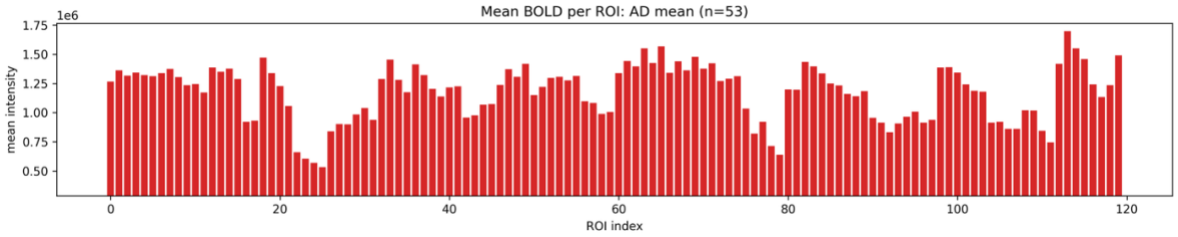

[A1]

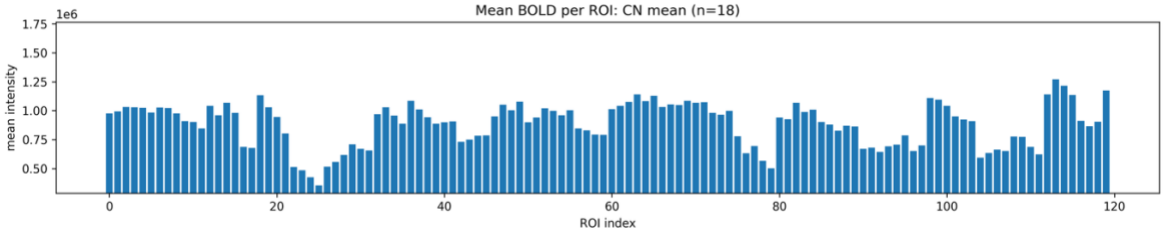

[A2]

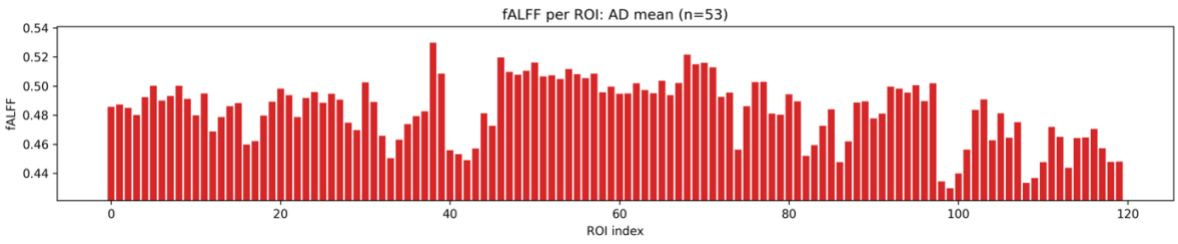

[A3]

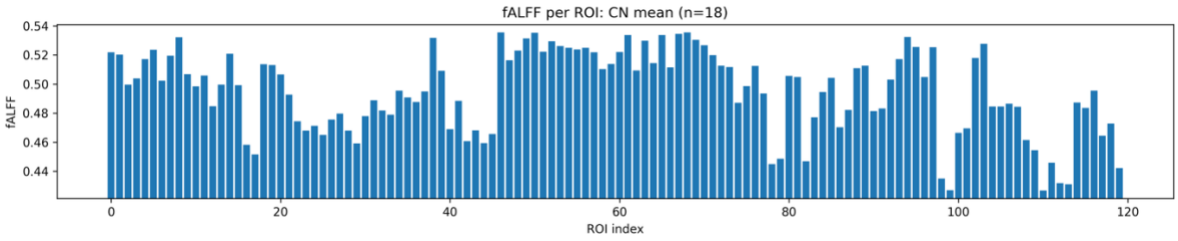

[A4]

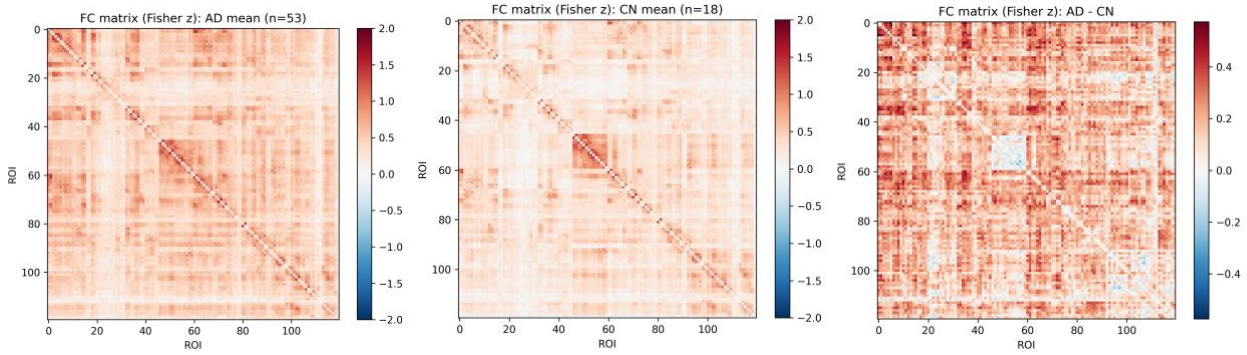
FC matrix (Fisher z): AD mean (n=53)
FC matrix (Fisher z): CN mean (n=18)
FC matrix (Fisher z): AD - CN
ROI
ROI

[A5] **[A6]** **[A7]**

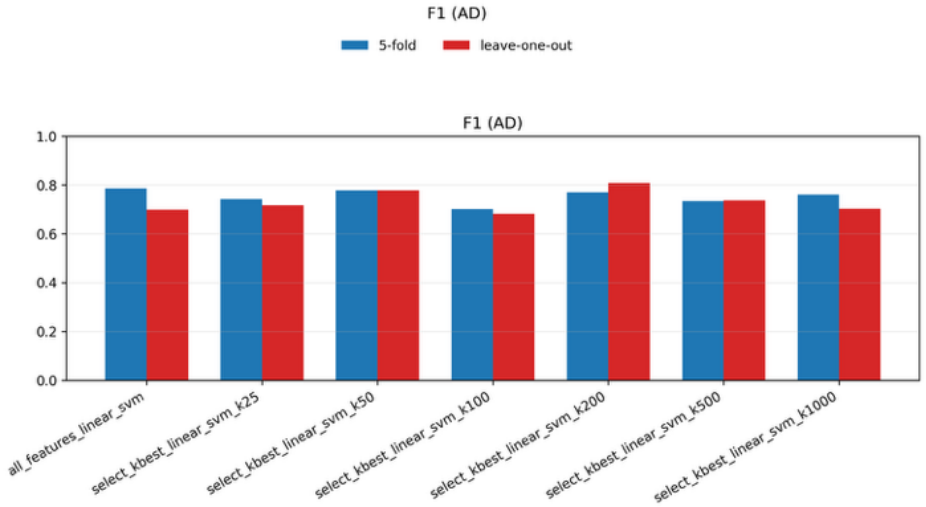
F1 (AD)
5-fold
leave-one-out
all_features_linear_svm
select_kbest_linear_svm_k25
select_kbest_linear_svm_k50
select_kbest_linear_svm_k100
select_kbest_linear_svm_k200
select_kbest_linear_svm_k500
select_kbest_linear_svm_k1000

[B1]

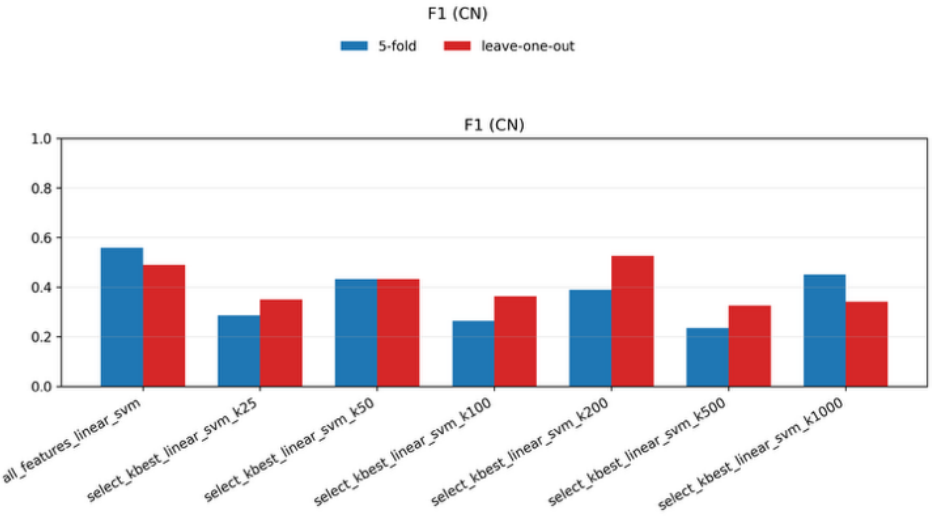
F1 (CN)
5-fold
leave-one-out
all_features_linear_svm
select_kbest_linear_svm_k25
select_kbest_linear_svm_k50
select_kbest_linear_svm_k100
select_kbest_linear_svm_k200
select_kbest_linear_svm_k500
select_kbest_linear_svm_k1000

[B2]

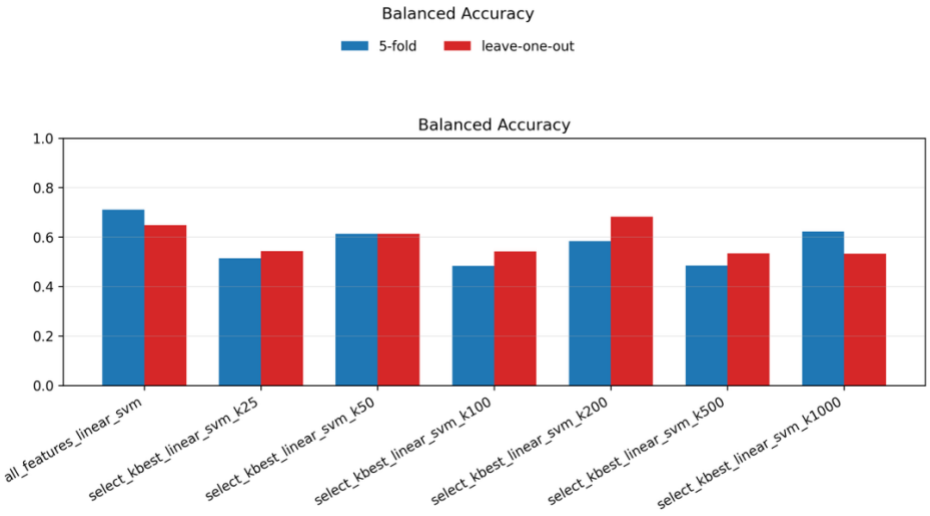

[B3]

References.

[1] Biswal, Bharat, et al. "Functional connectivity in the motor cortex of resting human brain using echo-planar MRI." *Magnetic resonance in medicine* 34.4 (1995): 537-541.

[2] Esteban, Oscar, et al. "fMRIPrep: a robust preprocessing pipeline for functional MRI." *Nature methods* 16.1 (2019): 111-116.

[3] Power, Jonathan D., et al. "Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion." *Neuroimage* 59.3 (2012): 2142-2154.

[4] Rolls, Edmund T., Marc Joliot, and Nathalie Tzourio-Mazoyer. "Implementation of a new parcellation of the orbitofrontal cortex in the automated anatomical labeling atlas." *Neuroimage* 122 (2015): 1-5.

[5] Fisher, Ronald A. "Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population." *Biometrika* 10.4 (1915): 507-521.

[6] Yu-Feng, Zang, et al. "Altered baseline brain activity in children with ADHD revealed by resting-state functional MRI." *Brain and Development* 29.2 (2007): 83-91.

[7] Zou, Qi-Hong, et al. "An improved approach to detection of amplitude of low-frequency fluctuation (ALFF) for resting-state fMRI: fractional ALFF." *Journal of neuroscience methods* 172.1 (2008): 137-141.

[8] Zang, Yufeng, et al. "Regional homogeneity approach to fMRI data analysis." *Neuroimage* 22.1 (2004): 394-400.

[9] Cortes, Corinna, and Vladimir Vapnik. "Support-vector networks." *Machine learning* 20.3 (1995): 273-297.

[10] Pedregosa, Fabian, et al. "Scikit-learn: Machine learning in Python." *the Journal of machine Learning research* 12 (2011): 2825-2830.

[11] Jack Jr, Clifford R., et al. "The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods." *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* 27.4 (2008): 685-691.