BRIEF REPORT

Functional Connectome Fingerprinting Using Shallow Feedforward Neural Networks

Gokce Sarar[a,b], Bhaskar Rao[b], and Thomas Liu[a,c,1]

[a]Center for Functional MRI, University of California San Diego; [b]Department of Electrical Engineering, University of California San Diego; [c]Departments of Radiology and Psychiatry, University of California San Diego

This manuscript was compiled on October 19, 2020

Although individual subjects can be identified with high accuracy using correlation matrices computed from resting-state functional magnetic resonance imaging (rsfMRI) data, the performance significantly degrades as the scan duration is decreased. Recurrent neural networks can achieve high accuracy with short duration (72s) data segments but are designed to use temporal features not present in the correlation matrices. Here we show that shallow feedforward neural networks that rely solely on the information in rsfMRI correlation matrices can achieve state-of-the-art identification accuracies ($\geq 99.5\%$) with data segments as short as 20s and across a range of input data-size combinations when the total number of data points (# regions $\times$ # time points) is on the order of $10{,}000$.

fMRI | Connectome | Neural Network | Fingerprinting

Functional connectome fingerprinting based on the similarity of correlation coefficient matrices computed from rsfMRI data can identify individuals with high accuracy ($> 98\%$) using long duration (> 12 minute) scans but considerably lower accuracy ($\approx 68\%$) is obtained when the data duration is decreased to 72s (1). Recurrent neural networks (RNN) can achieve high accuracy (98.5%) with short duration (72s) data, presumably reflecting their ability to capture both spatial and temporal features (2, 3). However, it has been shown that high RNN performance can be achieved even when the temporal order of the fMRI data is permuted (4), suggesting that the temporal features are not critical for identification. Here we introduce two shallow feedforward neural networks that can achieve high identification accuracy without the need for recurrent connections. Furthermore, we use these networks to estimate the minimum size of the data needed to robustly identify subjects with high mean accuracy ($\geq 99.5\%$) from short segments of rsfMRI data.

The two networks considered are shown in Figure 1A and 1B. The input to the correlation neural network (corrNN) consists of the upper triangular elements of the correlation coefficient matrix C estimated from a data matrix X consisting of z-normalized time series (of length N) from M regions of interest (ROI). For identification of L subjects, the network structure consists of a fully connected classification layer with L units, a batch normalization layer, and a softmax layer. The norm-based neural network (normNN) uses the z-normalized data X as the input. The first stage is a fully connected layer that projects the data onto K hidden units using the $M \times K$ weight matrix W to form the $N \times K$ intermediate matrix $Y = XW$. In the second stage, the L_2 norm across the time-dimension (i.e. across each column of Y) is computed for each hidden unit to form a summary measure of similarity over the collection of N time points. The resulting vector $F = \sqrt{diag(Y^T Y)} = \sqrt{diag(W^T C W)}$ is comprised of K features extracted from the correlation matrix C. The kth feature is proportional to the variance in the direction of the kth column vector of W. If these vectors are randomly oriented and constrained to be unit norm, then the features represent a random sampling of the "peanut" shaped surface of directional variances (5). The subsequent stages in the network are: a batch normalization layer, a fully connected classification layer with L hidden units, a second batch normalization layer, and a softmax layer.

Results

We assessed the performance of the two networks using data from the Human Connectome Project (HCP) (6). Two rsfMRI scans acquired on Day 1 were used for training, while the two scans from Day 2 were used for validation and testing.

For $M = 379$ ROIs, $N = 100$ time points (72s duration) per segment, and $K = 256$ hidden units, the mean classification accuracies of the corrNN and normNN models were 99.9% and 99.2%, respectively, for an initial set of 100 subjects, and 99.8% and 99.5% for a second independent set of 100 subjects. These accuracies are higher than those reported (94.3% to 98.5%) for RNN models (2, 3). For comparison, the mean classification accuracy using the similarity of the correlation coefficients was 79.4% for 100 time points per segment, which is higher than the 68% mean accuracy reported in (1) using data from a different dataset.

We used a greedy search algorithm to assess the relative importance of the ROIs with respect to model accuracy. Importance maps are shown in the top rows of Figures 1C and 1D, respectively, with the subsequent rows thresholded to highlight the top 15 to 60 ROIs. When considering the top 60 ROIs, the highest number of ROIs for both models are found in the dorsolateral prefrontal cortex followed by inferior parietal cortex, lateral temporal cortex, superior parietal cortex (for CorrNN), inferior frontal Cortex, and dorsal stream visual cortex, where brain regions are as defined in (7).

We used the top ROIs to evaluate CorrNN and NormNN performance with 15 to 60 ROIs and 5 to 1000 time points, as shown in Figures 1E and 1F, respectively. As the number of ROIs decreases, the number of time points needed to achieve higher accuracy increases. Defining 99.5% as the threshold for high mean accuracy, we observed that this threshold is surpassed with as few as $M = 60$ ROIs and $N = 100$ time points for CorrNN and 40 ROIs and 200 time points for NormNN,

G.S., B.R., and T.L. designed the research. G.S. and T.L. performed the research and wrote the paper.

The authors declare no competing interest.

[1]To whom correspondence should be addressed. E-mail: ttliu@ucsd.edu

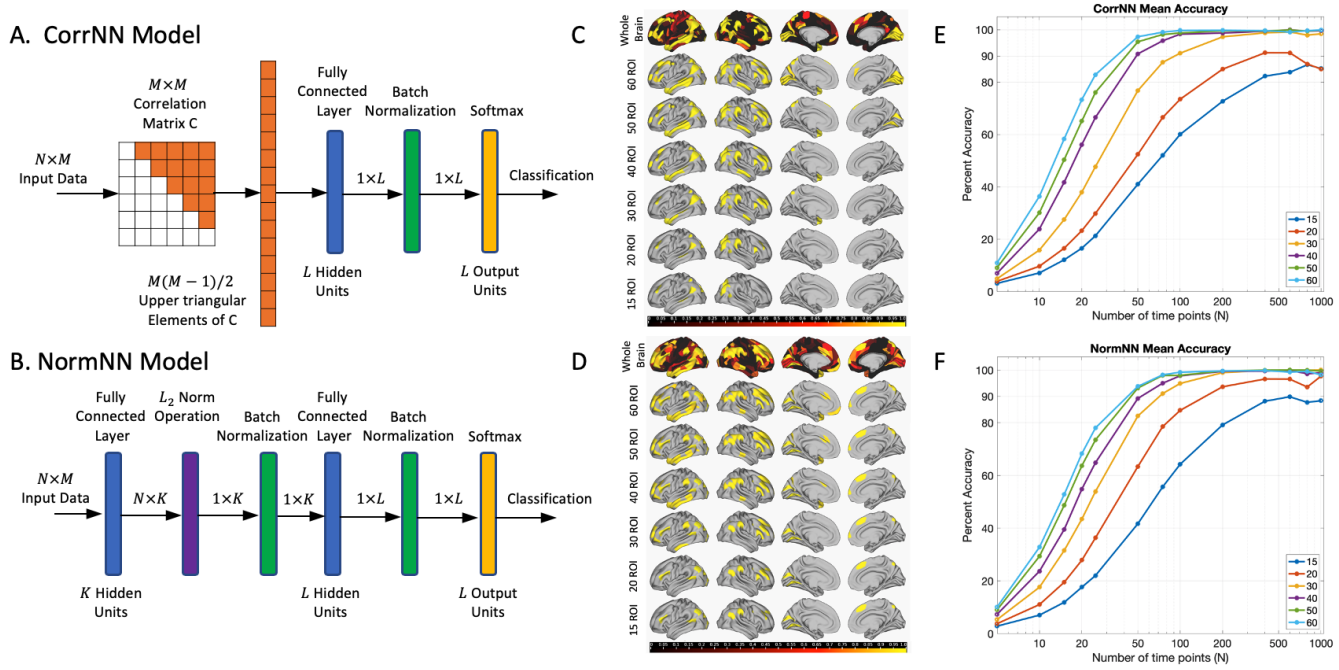

Fig. 1. (A,B) CorrNN and NormNN model structures. (C,D) Top rows: Maps showing the relative importance of the ROIs for identification accuracy with maximum importance of 1.0 indicated in yellow. The remaining rows are thresholded to show the locations of the top 15 to 60 ROIs. (E,F) Mean identification accuracies as a function of the number of time points and ROIs.

corresponding to $M \times N = 6000$ or 8000 total data points, respectively.

To further explore the dependence on number of ROIs and time points, we considered combinations (M, N) where the total number of data points was constrained to be equal to or close to either 6000 or 10,000 (see Figure 2 caption). For CorrNN, high mean accuracies are obtained for two of the combinations (dark red squares) with 6000 data points and all five of the combinations (dark red diamonds) with 10,000 points, respectively.

For NormNN, the number of parameters exhibits a linear dependence on the number of ROIs (M) as compared to the quadratic dependence for CorrNN (see Figure 2 caption). To better compare the models, we increased K by powers of 2 up to the value $K_{eq} = 0.5L\left(M^2 - M\right)/(M+L+3)$ for which the numbers of NormNN and CorrNN parameters were equivalent, while also including K_{eq} as one of the possible options. In figure 2b, we show NormNN accuracies obtained for either (1) the minimum value of $K \geq 256$ that surpassed the 99.5% threshold or (2) the value $K \leq K_{eq}$ that achieved the highest accuracy when the threshold was not met. High mean accuracies were obtained for three and four of the combinations with 6000 and 10,000 data points, respectively. As shown by the histograms, the high mean CorrNN and NormNN accuracies correspond to robust identification performance with the majority of the trials demonstrating 100 percent prediction accuracy.

Using the ROIs determined from the first 100 subjects we evaluated performance on the 2nd set of 100 subjects for the combinations denoted in Figure 2. High mean CorrNN accuracies ($\geq 99.5\%$) were maintained for both of the previously identified high performance combinations with 6000 points and three of the combinations with $10,000$ points, with the remaining two combinations (with $M \geq 300$) exhibiting slightly lower accuracies ($\geq 99.33\%$) for the 2nd dataset. Thus, the same set of ROIs can offer comparable and high levels of performance across independent datasets.

For NormNN we find that the first layer trained weights are randomly distributed so that the features after the L_2 norm operation represent an approximately uniform sampling of the directional variance surface of C. Indeed, high performance can also be achieved by replacing the first layer with a set of random Gaussian weights. The generalizability of the features across datasets exhibits a dependence on the number of units K. For example, when using first layer weights trained using the first set of subjects, performance for the combination $(50, 200)$ with $K = 256$ drops from 99.66% for the first 100 subjects to 97.98% for the second 100 subjects. Increasing to 512 units with weights trained using the first set yields accuracies of 99.64% and 99.52% for the first and second sets, respectively. Essentially the same accuracy levels (99.59% and 99.56%) are obtained when using random weights for the first layer. Thus, generalizability of the NormNN features increases when there is a higher number of features to characterize the directional variance.

Discussion

We have shown that shallow feedforward models can identify subjects based solely on information in rsfMRI correlation matrices, with CorrNN directly using the correlation coefficients as features while NormNN uses features related to the directional variance surface. The performance levels achieved are state-of-the-art, with high ($\geq 99.5\%$) mean identification accuracies robustly obtained with 6000 to 10,000 data points. For comparison, the convolutional RNN presented in (3) achieved 98.5% accuracy with $23,600$ data points.

Consistent with prior observations (1), high performance

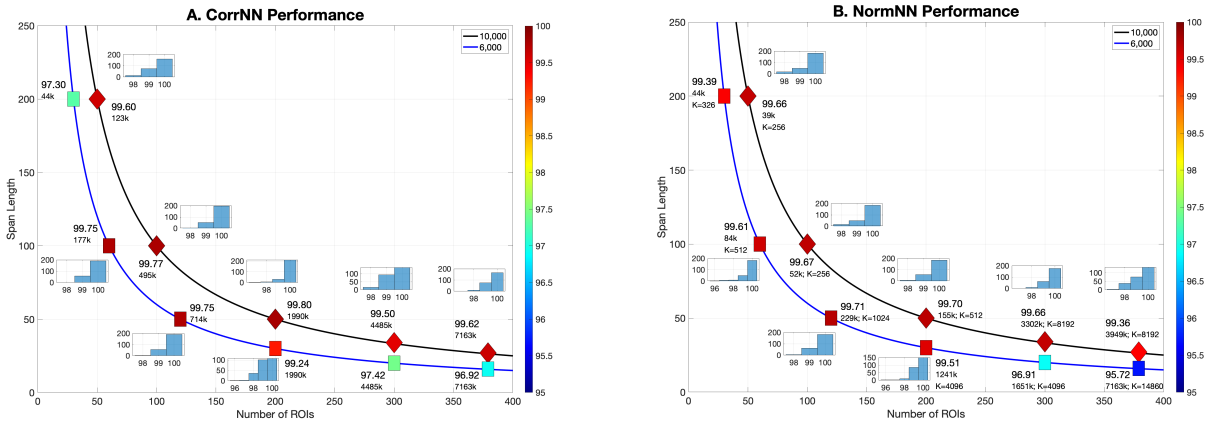

Fig. 2. (A,B) CorrNN and NormNN identification accuracies for combinations (M, N) of numbers of ROIs (M) and span lengths (N) that are constrained to have either 6000 or 10,000 data points, with the exception of the combinations $(379, 16)$, $(300, 34)$, and $(379, 27)$, which have 6064, 10,200, and 10,233 data points, respectively. Mean accuracies are indicated by labels and color scale. The numbers of model parameters (in thousands) for CorrNN $\left(0.5L\left(M^2 - M + 6\right)\right)$ and NormNN $(K(M + L + 3) + 3L)$ are also listed (with $L = 100$), as are the numbers of hidden units (K) for NormNN combinations. For combinations where CorrNN mean accuracy is greater than 99%, auto-scaled histograms show the distribution of identification accuracies obtained over 250 test trials per combination.

can be achieved when using a subset of the ROIs, including those located in frontoparietal and lateral temporal regions. The same set of ROIs can be used to achieve high performance across independent datasets, suggesting that the predictive value of inter-subject variability in the functional boundaries and connectivity of these regions generalizes across datasets.

While combinations with span lengths as short as 27 points (19.5s; CorrNN $(379, 27)$) can offer high performance, they require a large number of model parameters. In contrast, combinations with fewer ROIs but increased span lengths (e.g. $(100, 100)$) achieve high performance with one to two orders of magnitude fewer parameters. For NormNN the number of trainable parameters can be further decreased through the use of random weights in the first layer.

The effectiveness of the feedforward networks for distinguishing individuals with relatively little data suggests that similar future approaches may have the potential to more fully utilize the information contained in rsfMRI data to better identify disease-related differences.

Materials and Methods

Data and Preprocessing. We used pre-processed rsfMRI data from the 100 Unrelated Subjects subset of the HCP1200 Data Release and an additional 100 subjects from the S900 subset. HCP preprocessing included detrending, denoising, and registration to a common cortical surface. We divided the data into 379 regions of interest (ROI) consisting of 360 ROIs in the cortex as defined in (7) and 19 subcortical ROIs as defined in the group average parcellation provided by the HCP (6). The time series were averaged within each ROI and global signal regression was applied to the ROI time series.

Training, Validation, and Testing. We used the ROI-averaged data from day one as the training set (2 scans; 1200 pts per scan) and data from day two as validation and test sets (1 scan each). The performance was initially assessed using $K = 379$ ROIs and $N = 100$ time points. Overlapping data segments were used with shifts of 1 point between segments for training and 25 points for both validation and testing. The full connected layers were initialized with the Glorot uniform initializer We used the Adam optimizer with learning rate, β_1,and β_1 values of 0.001, 0.9, and 0.999, respectively. We used a batch size of 64 and monitored the validation loss every 600 steps with patience set to 30 monitoring steps, with learning rate annealed by halving it every 100 monitoring points.

Using the trained model weights, we assessed the relative importance of the ROIs by first looping over all 379 ROIs, independently zeroing out the data from each ROI, and finding the ROI for which the model retained its maximum accuracy. The identified ROI data was eliminated (i.e. set to zero) for the remainder of the process and the search was repeated over the remaining 378 ROIs to select the next ROI for elimination. This process was continued until only one ROI remained. The importance score of each ROI was $1 - M_c$ where M_c denotes the model accuracy just prior to ROI elimination.

We then examined performance across a range of ROI numbers ($K = 15$ to 60; with the top K ROIs selected based on importance scores) and durations ($N = 5$ to 1000 time points), retraining the models for each combination (K, N) with 1 point shifts for both training and validation. Testing was performed with 250 randomly chosen initial starting points with the exception of 200 sequential points for $N = 1000$. Performance using additional parameter combinations was also evaluated.

We used the approach of (1) to identify subjects based on similarity of the correlation coefficient matrices. Target matrices were calculated using all the data from day one, whereas test matrices were calculated using data segments from day two. Identification was performed by computing the spatial correlation between the test and target matrices and then matching (with replacement) each test matrix to the most highly correlated target matrix.

ACKNOWLEDGMENTS. We thank Eric Wong, Garrison Cottrell, Jiawei Ren, and Shili Wang for their assistance and helpful comments.

1. Finn ES, et al. (2015) Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. *Nature Neuroscience* 18(11):1664–1671.
2. Chen S, Hu X (2018) Individual Identification Using the Functional Brain Fingerprint Detected by the Recurrent Neural Network. *Brain Connectivity* 8(4):197–204.
3. Wang L, Li K, Chen X, Hu XP (2019) Application of Convolutional Recurrent Neural Network for Individual Recognition Based on Resting State fMRI Data. *Frontiers in Neuroscience* 13(May):1–8.
4. Sarar G, Wang S, Ren J, Liu TT (2019) Functional Connectome Fingerprinting using Recurrent Neural Networks Does Not Depend on the Temporal Structure of the Data in *Proceedings of the 27th Annual Meeting of the ISMRM.*
5. Bartz D, Hatrick K, Hesse CW, Müller KR, Lemm S (2013) Directional Variance Adjustment: Bias Reduction in Covariance Matrices Based on Factor Analysis with an Application to Portfolio Optimization. *PLoS ONE* 8(7):e67503.
6. Van Essen DC, et al. (2013) The WU-Minn Human Connectome Project: An overview. *NeuroImage* 80:62–79.
7. Glasser MF, et al. (2016) A multi-modal parcellation of human cerebral cortex. *Nature* 536(7615):171–178.