2021 43rd Annual International Conference of the
IEEE Engineering in Medicine & Biology Society (EMBC)
Oct 31 - Nov 4, 2021. Virtual Conference

Deep learning in resting-state fMRI*

Anees Abrol, Reihaneh Hassanzadeh, Sergey Plis and Vince Calhoun

Abstract— **Modeling the rich, dynamic spatiotemporal variations captured by human brain functional magnetic resonance imaging (fMRI) data is a complicated task. Analysis at the brain's regional and connection levels provides more straightforward biological interpretation for fMRI data and has been instrumental in characterizing the brain thus far. Here we hypothesize that spatiotemporal learning directly in the four-dimensional (4D) fMRI voxel-time space could result in enhanced discriminative brain representations compared to widely used, pre-engineered fMRI temporal transformations, and brain regional and connection-level fMRI features. Motivated by this, we extend our recently reported structural MRI (sMRI) deep learning (DL) pipeline to additionally capture temporal variations, training the proposed 4D DL model end-to-end on preprocessed fMRI data. Results validate that the complex non-linear functions of the used deep spatiotemporal approach generate discriminative encodings for the studied learning task, outperforming both standard machine learning (SML) and DL methods on the widely used fMRI voxel/region/connection features, except the relatively simplistic measure of central tendency - the temporal mean of the fMRI data. Additionally, we identify the fMRI features for which DL significantly outperformed SML methods for voxel-level fMRI features. Overall, our results support the efficiency and potential of DL models trainable at the voxel level fMRI data and highlight the importance of developing auxiliary tools to facilitate interpretation of such flexible models.**

I. Introduction

Brain function is assumed to be spatially segregated, i.e., different aspects of function tend to be linked to different brain sites or different combinations of brain networks. Blood oxygenation level dependent (BOLD) functional magnetic resonance imaging (fMRI) assesses brain function by capturing regional, temporal changes due to the hemodynamic and metabolic sequelae to underlying elevated neuronal activity. To understand brain function via the fMRI modality, it is vitally important to thoroughly evaluate the variations captured by the data over both space and time.

Analysis at the brain's regional and connection levels facilitates biological interpretation of fMRI data and has been instrumental in characterizing the brain thus far. Previous work in the field has typically reduced the fMRI data's dimensionality in several unique ways in space and time. Region of interest (ROI) or atlas-based approaches, as well as data-driven parcellation methods such as group independent component analysis (gICA), are often used to reduce the fMRI data in space and discover task-discriminative information at the brain region level. Such decomposition allows identifying a set of brain components (i.e., spatial clusters of brain voxels) and associated activity time-courses (ICA-TCs). Subsequently, simplistic pairwise statistical measures of synchronous activation/de-activation (e.g., Pearson's correlation, coherence, mutual information, etc.), also termed as functional connectivity (FC) measures, can be assessed for the estimated ICA-TCs at a static (sFC) or time-varying/dynamic (dFC) scale. Likewise, several temporal transformations such as the amplitude of low-frequency fluctuations (ALFF), the fractional amplitude of low-frequency fluctuations (fALFF), Kendall's coefficient of concordance Regional Homogeneity (KccReHo), voxel-mirrored homotopic connectivity (VMHC), degree centrality (DC), Percent Amplitude of fluctuation (PerAF), etc. have been proposed to reduce the data in time while preserving the voxel dimensionality.

Compared to the above-mentioned valuable fMRI features, modeling the rich, dynamic spatiotemporal variations to identify biomarkers directly from the four-dimensional (4D) fMRI space is perhaps an even more complicated, computationally demanding task. One potential way to explore this objective is by leveraging the flexibility of deep learning (DL) models, as these have already shown great promise in diverse medical imaging and genomic data applications over the last few years [1-6]. However, most DL brain imaging studies work with pre-engineered features (i.e., summary measures region or connection level), thus depriving DL of its fundamental advantage – representation learning from raw or minimally preprocessed voxel (i.e., input) space.

Our recent work [7] demonstrates, on multiple learning tasks on structural MRI (sMRI) data, that if trained following prevalent DL practices, DL methods have the potential to scale particularly well and substantially improve compared to SML methods. Specifically, we observed that our three-dimensional (3D) CNN-based DL model, trained end-to-end on gray matter maps (i.e., directly in the voxel (3D) space), generated superior brain representations to outperform standard machine learning (SML). Following that, here we hypothesize the same may hold true for BOLD fMRI data – that spatiotemporal learning directly in the 4D fMRI voxel-space could discover enriched discriminative brain representations compared to the widely used, pre-engineered fMRI temporal transformations in 3D space or ICA-based features.

To test the feasibility and efficacy of spatiotemporal learning on fMRI data, we develop a sample 4D DL model (referred to as the 4DStL model) by extending our previous DL pipeline [7] with an extra convolutional layer upfront to capture temporal variation in an end-to-end manner. Specifically, we use 1x1x1 kernels in the added convolutional

* Research supported by National Institutes of Health (NIH) grants R01EB006841, RF1AG063153, and R01EB020407 to Dr. Vince D. Calhoun.

A.A., R.H., S.P. and V.C. are with the Tri-institutional Center for Translational Research in Neuroimaging and Data Science (TReNDS), Georgia State University, Georgia Institute of Technology, Emory University, Atlanta, GA 30303 USA. (Corresponding author: Anees Abrol, e-mail: abrolanees@gmail.com).

layer to identify temporal descriptors without changing the following convolutional layers' receptive field. This step reduces the timepoint dimension in the input tensor to a predefined number of outputs channels, while end-to-end training upholds optimal weight estimation. We profile this comparison on the age regression learning task, wherein the objective is to minimize the mean absolute error (MAE) between the actual and the predicted age, a popular proxy benchmark for evaluating other scientifically captivating questions, e.g., cognitive function, mental disorders. Figure 1 describes the overview of the study design.

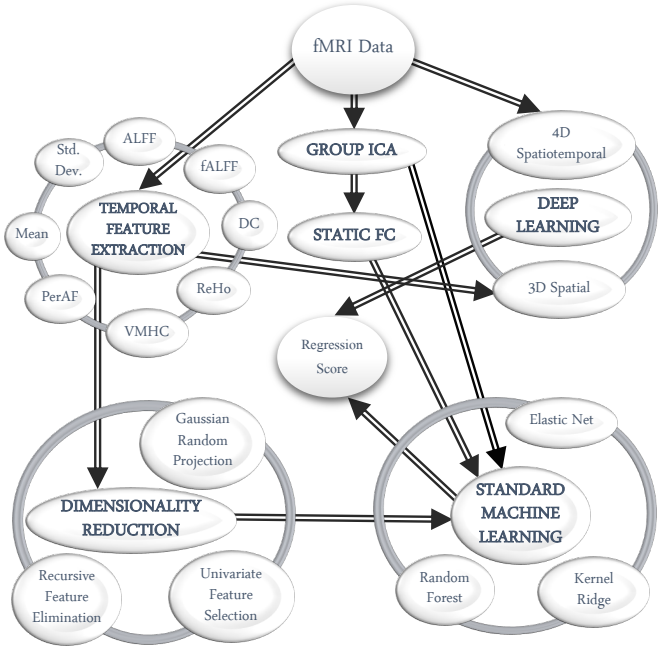

Figure 1: Study Workflow. The figure describes the overview of the study design. This work compares the performance of (a) several temporal transformations on fMRI data, group ICA time-courses, and static functional connectivity for a sample learning/regression task using a range of SML and DL methods.

II. Materials and Methods

A. Data

This benchmark was established on fMRI images (n = 12,314) from unaffected subjects (i.e., those who had no diagnosed or self-reported mental illnesses based on 22,392 subjects' fMRI data available as of 7 April 2019) from the UK Biobank repository. The scientific study protocol of the UK Biobank is approved by the Ethics and Governance Council. Written informed consent was obtained from all subjects participating in the UK Biobank study.

The UK Biobank study [8] used a Siemens Skyra 3T running VD12A SP4 scanner with a standard Siemens 32-channel RF receive coil. The gradient-echo echo-planar imaging (GE-EPI) data parameters included - slice acquisition: multislice (multiband) acceleration, phase-encoding direction: AP (anterior-posterior), resolution: 2.4x2.4x2.4 mm, field-of-view: 88x88x64 matrix, duration: 6 minutes with a TR of 0.735 s and TE of 39 ms (a total of 490 timepoints), and GE-EPI: x8 multislice acceleration, no iPAT, flip angle: 52°.

B. fMRI Data Preprocessing

The UK Biobank fMRI data minimally preprocessed via the Melodic pipeline [8] is used in this work. This preprocessing pipeline featured motion-correction using MC-FLIRT, grand-mean intensity normalization of the entire 4D dataset by a single multiplicative factor, high-pass temporal filtering (Gaussian-weighted least-squares straight-line fitting, with sigma = 50.0s), EPI and GDC unwarping, and removal of structured artefacts by ICA+FIX processing. The data was further registered to the MNI EPI template using FMRIB's Linear Image Registration Tool (FLIRT) followed by normalization to the SMP12 old normalization module and smoothed using a Gaussian kernel with FWHM = 6mm. The preprocessed fMRI maps used for each of the 12314 subjects included 490 volumes of 53 x 63 x 52 voxels.

C. Comparative fMRI Feature Extraction

We estimated several temporal transformations on fMRI data such as ICA-TC, sFC, ALFF, fALFF, KccReho, VMHC, DC, and PerAF in this study, moderately motivated by previous work [7, 9, 10]. The preprocessed fMRI data were decomposed using spatially-constrained group ICA using the Neuromark template as reference maps [11] (n=53 brain regions) to estimate the group-ICA time-courses (ICA-TCs). Subsequently, the sFC features ($n={}^{53}C_2=1378$ brain connections) were computed as pairwise Pearson's correlation coefficients (unthresholded) between the ICA-TCs. The performance of all fMRI features was measured by the SML and DL methods outlined in Section II-E. Additionally, relatively simplistic measures of central tendency (e.g., the temporal mean of fMRI data – TMF) and dispersion (e.g., temporal standard deviation – TSF) were also computed, and their combination stacked in the channel dimension of the input tensor (TMSF) was also evaluated motivated by previous work [12]. Feature extraction is an indispensable prerequisite to eliminate redundant features and boost the performance of the SML methods. Therefore, similar to our previous work [12], we reduced all of these voxel-level features (i.e., all except sFC) to a lower dimensionality (n=1000) with three dimensionality reduction methods - Gaussian random projection (GRP), recursive feature elimination (RFE), and univariate feature selection (UFS).

D. Spatiotemporal DL Model

Notably, experimenting with the fast TR (0.735s) whole-brain UK Biobank data corresponds to training a high number of 3D volumes (n = 490), thus constituting a highly computationally demanding task even for a moderate batch size (e.g., n = 16 subjects). To alleviate this problem, we smoothed the fMRI dataset using piecewise aggregate approximation (PAA) of the fMRI time series with a window size of 15 time points, thus effectively reducing the data temporally (from n = 490 to n = 32 time samples), but still retaining appreciable temporal information. We train the 4DStL model on this smoothed data to capture the temporal variations using 1x1x1 kernels in its foremost convolutional layer. This step allowed us to identify temporal descriptors to a predefined number of output channels (n=3) without changing the following convolutional layers' receptive field, thus allowing reuse of our proven 3D DL model [7] for spatiotemporal learning via the proposed 4DStL model. Therefore, this model employs a 3C-64C-128C-192C-192C-

128C configuration of convolutional layers, wherein the numbers signify the number of channels/filters for a given convolutional layer.

E. Comparative SML and 3D DL Models

We evaluated all fMRI features using three SML models - elastic-net (EN), kernel ridge regression (KRR), and random forest (RF) ensemble learning, as well as a suitable DL model to establish comparative baselines for the 4DStL model. Specifically, we used our recent 3D-CNN DL model [7] for learning all voxel-level temporal transformations, a two-layer bi-directional long short-term memory (LSTM) model [13, 14] for ICA-TCs and a graph-CNN model [15, 16] for sFC features (modified version of BrainNetCNN implementation in [9]). The three SML and DL models were implemented with the scikit-learn python machine learning library using our custom code [17]. Furthermore, the 3D DL model used in this work is a 3D-CNN variant of the AlexNet architecture [18] (64C-128C-192C-192C-128C) adapted to a deep vanilla regressor [17] by reducing the output nodes in the final fully connect layer to one.

F. Cross-validation and Hyperparameter-tuning

All analysis in this work was evaluated in a standard, repeated (n = 5), stratified cross-validation (CV) procedure and using the same training/validation/test data partitions. Hyperparameter tuning was employed for SML models through a grid parameter search with the hypopt python package. For the KRR method, the hyperparameter grids were spanned for the kernel mapping function (linear/radial-basis-function/polynomial/sigmoidal), regularization strength (alpha in the [10^{-3}, 10^{-2}, 10^{-1}, 1] range), and the gamma parameter in kernel mapping functions (10 values sampled on a logarithm scale for a range of powers of 10 from [−4, 2]). For the RF method, the number of trees in the forest (5 values sampled uniformly in the [100, 200] range), the maximum number of features considered at each split (chosen as square root or logarithm to the base 2 of the number of features), the minimum number of samples required to split an internal node (2, 5 or 10), the minimum number of samples required to be at a leaf node (1, 2, or 4) and the bootstrap flag (on/off) were tuned as hyperparameters. For the EN regression method, the alpha tuning parameter multiplied to the penalty terms (in the [10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}] range) and convex combination penalty parameter (10 values sampled uniformly in the [0,1] range) were tuned as hyperparameters.

Furthermore, for the 3D CNN DL model, we used a batch size of 16 and retained the learning rate (10^{-2}) as validated in the [10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}] range from the hyperparameter tuning stage of our previous work. A learning rate scheduler callback was employed to reduce the learning rate by a factor of 0.5 on plateauing of the validation accuracy metric, and the MSE loss was used. Early stopping with a patience level of 40 epochs was implemented on the validation MAE measure to reduce overfitting and achieve lower generalization error in the testing phase. The MAE metric was chosen for early stopping to optimize test MAE, thus allowing for a direct comparison to previous work. The 2-layer bidirectional LSTM model and graph-CNN models used the Adam optimizer and tuned for the learning rate (in the [5x10^{-2}, 10^{-2}, 5x10^{-3}, 10^{-3}, 10^{-4}] range) and batch size (16, 32, 64, 128). Additionally, for the LSTM model, the hidden layer size parameters were tuned (32, 64, 128, 256). Training and testing routines for the DL architectures were implemented on an NVIDIA CUDA parallel computing platform using GPU accelerated NVIDIA CUDA toolkit (cudatoolkit), CUDA Deep Neural Network (cudnn), and Pytorch tensor libraries.

III. Results

A. Performance of DL methods

The MAE, correlation and coefficient of determination (R2) regression scores between the actual and the predicted age were estimated for all features and models. Figure 2 illustrates the advantage of DL as compared to SML on the age regression task on fMRI data. The 4D spatiotemporal learning model (4DStL) reported a MAE of 3.54 (years), correlation (r) value of 0.81 and R2 value of 0.65, a significant improvement over ICA-TC, sFC and all temporal transformations except the temporal mean of fMRI data (TMF: MAE = 3.48, r = 0.82, R2 = 0.66), and the case when this measure was stacked with the temporal standard deviation (TMSF: MAE = 3.45, r = 0.82, R2 = 0.66). Additionally, the 3D DL model resulted in the least MAE for all other temporal transformations (MAE_{ALFF} = 4.36, MAE_{fALFF} = 4.52, MAE_{DC} = 5.00, $MAE_{KccReHo}$ = 4.58, MAE_{VMHC} = 4.78, MAE_{TSF} = 4.45, MAE_{VMHC} = 4.78) as well as ICA-TC ($MAE_{ICA\text{-}TC}$ = 4.66) and sFC features (MAE_{SFC} = 4.8), as compared to the performance of all three SML methods for these features respectively. Notably, the performance difference between SML and DL methods was significant ($p < 0.005$) only for the TMF, TSF, TMSF, ALFF, fALFF, PerAF and ICA-TC features.

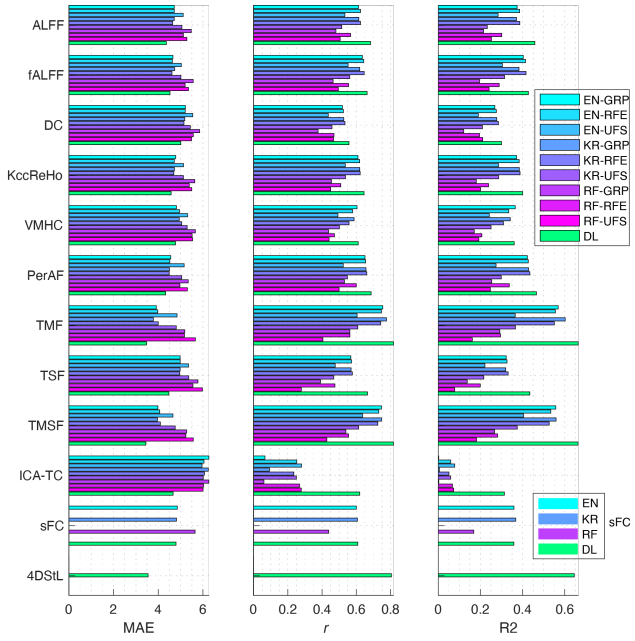

Figure 2: Performance comparison on the age regression task. Spatiotemporal deep learning was implemented using the 4DStL method directly on the 4D fMRI data. To establish a comparative baseline, different temporal transformations and features (ALFF, fALFF, DC, KccReHo, VMHC, PerAF, TMF, TMS, TMSF, ICA-TC and sFC) were computed from the preprocessed fMRI data. The performance of all methods and features for the entire range of SML and DL methods was cross-validated using the MAE, correlation (r) and coefficient of determination (R^2) regression metrics. The 4DStL model

on preprocessed fMRI data and 3D DL model on the TMF feature resulted in the smallest MAE on held-out test data.

B. *Performance trends in SML models*

For all fMRI features tested with SML methods, TMF reported the best performance when using the KR SML method for GRP features (MAE = 3.78 years, r = 0.78, R2 = 0.60), followed by the EN SML method for GRP features for the same temporal transformation (MAE = 3.91 years, r = 0.75, R2 = 0.57). Overall, when comparing the performance of the SML methods for all temporal transformations, EN and KR performed consistently the best for all the GRP and RFE features across all metrics, thereby also suggesting lower efficacy and value of UFS for dimensionality reduction and the RF method for regression tasks on such data.

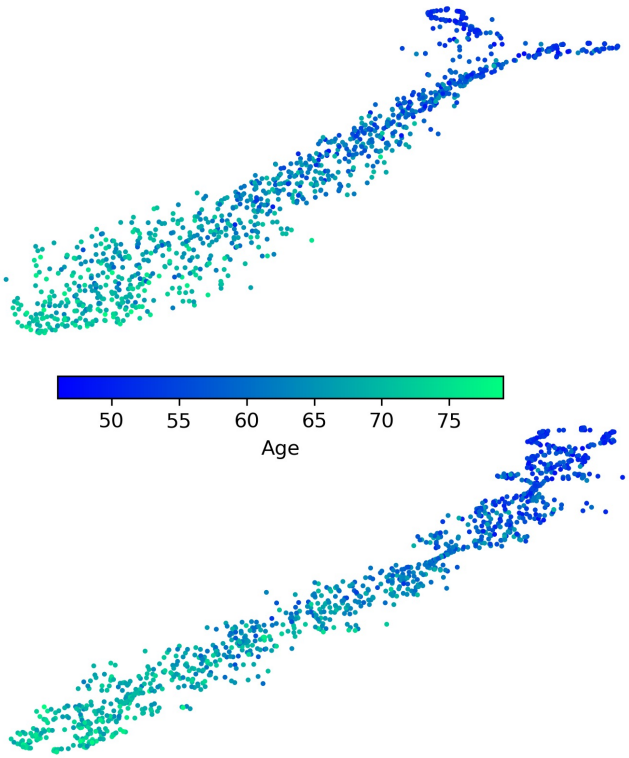

Figure 3: Non-linear embeddings of the representations encoded by the 4D 4DStL (deep spatiotemporal) model on preprocessed fMRI data (Top) and the 3D DL model on TMF features (Bottom). The task-discriminative fMRI data representations encoded by both models span comprehensible projection spectra.

C. *4DStL Model: Embeddings & Network Weights*

If the DL models indeed learn representations that encode the brain rationally, the encodings in the deeper layers (e.g., the last convolutional layer) must illustrate comprehensible projection spectra per the undertaken learning task. Figure 3 illustrates the non-linear embeddings of the fMRI data representations encoded by the 4DStL model on preprocessed fMRI data (top panel) and the 3D DL model on TMF features (bottom panel) as estimated using the *t*-distributed stochastic neighbor embedding (*t*SNE) framework. Evidently, subjects are seen ordered in increasing age from one end of the spectrum to the other except for some outliers, thereby confirming that these task-discriminative representations are consistently meaningful for both best-performing cases.

Additionally, Figure 4 introspects the filter weights of the first convolutional layer of the 4D spatiotemporal DL model. In this figure, the columns (rows) represent the 32 (3) input (output) channels. The non-uniform distribution of filter weights across the different output time channels of this model implies that the model was not trivially trained and that this pipeline may potentially discover temporal information presented by the fMRI data.

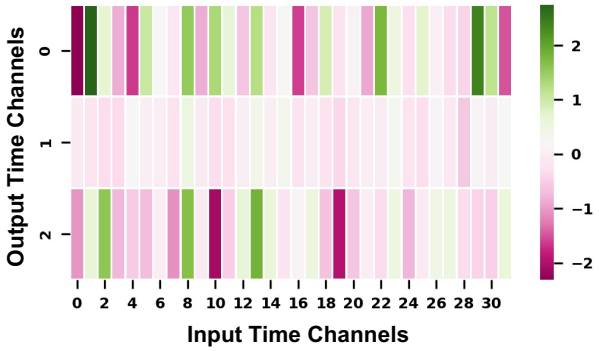

Figure 4: Distinct filter weights are visualized for the first convolutional layer of the 4DStL (deep spatiotemporal) model. This non-uniform distribution across the time channels suggests the model can discover temporal information in the fMRI data.

IV. Discussion

Our work systematically compares the performance of a wide range of temporal fMRI features evaluated using several SML and DL methods to an end-to-end deep spatiotemporal learning model directly on 4D fMRI data. We found that the used 4DStL model significantly outperformed all SML and DL methods compared to fMRI features except the fMRI data's temporal mean (TMF). This observation suggests that the undertaken 4DStL approach can learn vital spatiotemporal information and critically preserves voxel-wise interrelationships to an equitable extent similar to TMF features. Critically, the matched performance of TMF features doesn't necessarily imply a lack of meaningful temporal information at voxel level in the fMRI modality as there are several ways the tested 4DStL model may be tuned and perhaps numerous other flexible models that could be adapted to learn the temporal information in the fMRI data as discussed in the next paragraph. Instead, this particular observation further confirms the dominance of spatial/anatomical interrelationships in driving the undertaken learning task (i.e., age regression) and only suggests retaining that information may be more critical than the temporal dimension of this data on this learning task.

We also note that the used 4DStL model is potentially limited in a few ways, including (a) dimensionality reduction of the fMRI data in the time dimension (i.e., from 490 to 32) via the PAA approach, (b) reducing the data time dimension abruptly (i.e., from 32 to 3) using a 1x1x1 kernel in the first 3D convolutional layer and (c) deploying only one such layer as compared to reducing the data in time gradually with the use of multiple such layers. Exploring the above-mentioned potential limitations and confirming the trends suggested by this study on other learning tasks comprise exciting topics for future work in the domain. Nevertheless, it is evident from the

current results that the 4DStL model may also allow introspection in the time-domain, an attribute that is fundamental for explaining the fMRI data at all data stratum (i.e., the voxel, region, and connection levels).

Finally, while the 4DStL framework tested for this exploratory analysis is based on a CNN model, which is inarguably a more established and supervised DL model, investigating other flexible supervised/semi-supervised/unsupervised DL models, or perhaps rational combinations like CNN and an RNN (e.g., to learn the spatial and temporal properties of the data sequentially in an end-to-end manner) could be similarly interesting. We presume research in this direction should target the identification of deep spatiotemporal models that source a vibrant blend of excellent representational learning and precise biological interpretability for fMRI data.

V. Conclusion

The results in this work support the efficiency and potential of DL models trainable end-to-end on preprocessed fMRI data and highlight the importance of developing auxiliary tools to explain such flexible models. Overall, these observations motivate future work in the field to create and establish flexible DL models for spatiotemporal learning directly in the 4D space, ultimately aimed toward facilitating rich characterizations of the fMRI modality to explain vital neuroimaging objectives at all data stratum (i.e., the voxel, region, and connection levels).

Acknowledgment

The authors acknowledge and thank the TReNDS Center data preprocessing team – Dr. Rogers F. Silva, Dr. Yuhui Du, Mustafa Salman, Dr. Anna Bonkhoff, Dr. Zening Fu, Dr. Thomas DeRamus, Eswar Damaraju, Bradley Baker, and Dr. Anees Abrol, for preprocessing and quality control of the UK Biobank data used in this work.

References

[1] D. Shen, G. Wu, and H.-I. Suk, "Deep Learning in Medical Image Analysis," (in eng), *Annual review of biomedical engineering,* vol. 19, pp. 221-248, 2017, doi: 10.1146/annurev-bioeng-071516-044442.

[2] G. Hinton, "Deep Learning-A Technology With the Potential to Transform Health Care," (in eng), *Jama,* vol. 320, no. 11, pp. 1101-1102, Sep 18 2018, doi: 10.1001/jama.2018.11100.

[3] H. Greenspan, B. v. Ginneken, and R. M. Summers, "Guest Editorial Deep Learning in Medical Imaging: Overview and Future Promise of an Exciting New Technique," *IEEE Transactions on Medical Imaging,* vol. 35, no. 5, pp. 1153-1159, 2016, doi: 10.1109/TMI.2016.2553401.

[4] C. Angermueller, T. Pärnamaa, L. Parts, and O. Stegle, "Deep learning for computational biology," *Molecular Systems Biology,* https://doi.org/10.15252/msb.20156651 vol. 12, no. 7, p. 878, 2016/07/01 2016, doi: https://doi.org/10.15252/msb.20156651.

[5] G. Eraslan, Ž. Avsec, J. Gagneur, and F. J. Theis, "Deep learning: new computational modelling techniques for genomics," *Nature Reviews Genetics,* vol. 20, no. 7, pp. 389-403, 2019/07/01 2019, doi: 10.1038/s41576-019-0122-6.

[6] A. Abrol, M. Bhattarai, A. Fedorov, Y. Du, S. Plis, and V. Calhoun, "Deep Residual Learning for Neuroimaging: An application to Predict Progression to Alzheimer's Disease," *Journal of Neuroscience Methods,* p. 108701, 2020/04/08/ 2020, doi: https://doi.org/10.1016/j.jneumeth.2020.108701.

[7] A. Abrol *et al.*, "Deep learning encodes robust discriminative neuroimaging representations to outperform standard machine learning," *Nature Communications,* vol. 12, no. 1, p. 353, 2021/01/13 2021, doi: 10.1038/s41467-020-20655-6.

[8] K. L. Miller *et al.*, "Multimodal population brain imaging in the UK Biobank prospective epidemiological study," (in eng), *Nat Neurosci,* vol. 19, no. 11, pp. 1523-1536, Nov 2016, doi: 10.1038/nn.4393.

[9] T. He *et al.*, "Deep neural networks and kernel regression achieve comparable accuracies for functional connectivity prediction of behavior and demographics," *NeuroImage,* vol. 206, p. 116276, 2020/02/01/ 2020, doi: https://doi.org/10.1016/j.neuroimage.2019.116276.

[10] R. M. Thomas, S. Gallo, L. Cerliani, P. Zhutovsky, A. El-Gazzar, and G. van Wingen, "Classifying Autism Spectrum Disorder Using the Temporal Statistics of Resting-State Functional MRI Data With 3D Convolutional Neural Networks," (in English), *Frontiers in Psychiatry,* Original Research vol. 11, no. 440, 2020-May-15 2020, doi: 10.3389/fpsyt.2020.00440.

[11] Y. Du *et al.*, "NeuroMark: an automated and adaptive ICA based pipeline to identify reproducible fMRI markers of brain disorders," *NeuroImage: Clinical,* p. 102375, 2020/08/11/ 2020, doi: https://doi.org/10.1016/j.nicl.2020.102375.

[12] X. Li *et al.*, "2-Channel convolutional 3D deep neural network (2CC3D) for fMRI analysis: ASD classification and feature learning," in *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, 4-7 April 2018 2018, pp. 1252-1255, doi: 10.1109/ISBI.2018.8363798.

[13] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Processing,* vol. 45, no. 11, pp. 2673-2681, 1997, doi: 10.1109/78.650093.

[14] S. Hochreiter and J. Schmidhuber, "Long Short-term Memory," *Neural computation,* vol. 9, pp. 1735-80, 12/01 1997, doi: 10.1162/neco.1997.9.8.1735.

[15] J. Bruna, W. Zaremba, A. Szlam, and Y. Lecun, "Spectral Networks and Locally Connected Networks on Graphs," 12/20 2013.

[16] J. Kawahara *et al.*, "BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment," (in eng), *Neuroimage,* vol. 146, pp. 1038-1049, Feb 1 2017, doi: 10.1016/j.neuroimage.2016.09.046.

[17] A. Abrol, "aabrol, aabrol/SMLvsDL v1.0 (Version v1.0)." *Zenodo. http://doi.org/10.5281/zenodo.4309677,* 2020, December 7.

[18] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Commun. ACM,* vol. 60, no. 6, pp. 84–90, 2017, doi: 10.1145/3065386.