

Attend and Decode: 4D fMRI Task State Decoding Using Attention Models

Sam Nguyen

Brenda Ng

Alan D. Kaplan

Priyadip Ray

Lawrence Livermore National Laboratory, Livermore, CA, USA

SAMIAM@LLNL.GOV

NG30@LLNL.GOV

KAPLAN7@LLNL.GOV

RAY34@LLNL.GOV

Editors: Emily Alsentzer[⊗], Matthew B. A. McDermott[⊗], Fabian Falck, Suproteem K. Sarkar, Subhrajit Roy[‡], Stephanie L. Hyland[‡]

Abstract

Functional magnetic resonance imaging (fMRI) is a neuroimaging modality that captures the blood oxygen level in a subject’s brain while the subject either rests or performs a variety of functional tasks under different conditions. Given fMRI data, the problem of inferring the task, known as *task state decoding*, is challenging due to the high dimensionality (hundreds of million sampling points per datum) and complex spatio-temporal blood flow patterns inherent in the data. In this work, we propose to tackle the fMRI task state decoding problem by casting it as a 4D spatio-temporal classification problem. We present a novel architecture called *Brain Attend and Decode* (BAnD), that uses residual convolutional neural networks for spatial feature extraction and self-attention mechanisms for temporal modeling. We achieve significant performance gain compared to previous works on a 7-task benchmark from the large-scale Human Connectome Project-Young Adult (HCP-YA) dataset. We also investigate the transferability of BAnD’s extracted features on unseen HCP tasks, either by freezing the spatial feature extraction layers and retraining the temporal model, or finetuning the entire model. The pre-trained features from BAnD are use-

ful on similar tasks while finetuning them yields competitive results on unseen tasks/conditions.

Keywords: Neuroimaging, fMRI, task state decoding, attention models

1. Introduction

Functional MRI is a neuroimaging modality that measures spatial differences in blood oxygenation over time. As opposed to forming a static structural image of the brain, fMRI captures temporal activity patterns. A contrast mechanism is used to generate the Blood Oxygen Level Dependent (BOLD) signal which may be used to infer brain activities by measuring relative oxygen concentrations in the blood [Webb and Kagadis \(2003\)](#). During the fMRI experiment, a subject lies in an MRI scanner while performing various tasks to activate different brain regions. The goal of fMRI task state decoding is to map the fMRI sequence to the task that was performed.

Functional MRI modality can help clinicians diagnose neurological diseases as well as monitor the effectiveness of therapy. For example, resting state fMRI identified significant disruptions in default mode networks coactivation in patients with Alzheimer’s disease [Greicius et al. \(2004\)](#); [Rombouts et al.](#)

(2009). In the case of Parkinson disease, which is a neurological movement disorder, Tada (1998); Sabatini et al. (2000) identified lowered activation in the supplementary motor area during movement in Parkinson versus controls. These signatures can potentially be used to diagnose neurological disorders in a noninvasive manner.

Our paper’s main focus is on the problem of fMRI task state decoding. Our attention-based model is geared towards the representation of complex spatio-temporal signals that fMRI data are comprised of. Beyond the formulation and validation of the models that we present in this paper, it may be possible to use them to (1) analyze the differences in fMRI data from patients versus controls, (2) directly classify neurological disorders, or (3) distinguish various pathological subtypes of a neurological disease to monitor impact of therapy specific to that subtype.

While many analytic techniques, such as Naselaris et al. (2011); Wen et al. (2018), have been developed for fMRI task state decoding, many challenges remain. Modern scanners are constantly improving its acquisition rates (number of frames per second) and its image resolution (number of pixels per frame). The data takes up significant memory and thus requires scalable solutions. Moreover, the phenomenon being measured – the blood oxygenation levels within the brain – is of high spatio-temporal complexity. To model this phenomenon, we need an expressive enough model that can capture this complexity. In this paper, we “leave no stone unturned” by using as much information as possible that is inherent in the raw fMRI data. We explicitly account for temporal dynamics in our model, and use it to compute spatio-temporal maps to indicate the relative importance of brain regions for the task decoding.

Our proposed model, called *Brain Attend and Decode* (BAnd), has been validated us-

ing the Human Connectome Project (HCP) Van Essen et al. (2013); Glasser et al. (2016b) neuroimaging dataset, one of the most prevalent datasets for brain imaging studies. HCP contains multiple types of neuroimaging, demographic information, and cognitive scoring of over 1,000 healthy young adults (HCP-YA). Our approach operates on volumetric fMRI data and does not require any computationally expensive parcellation techniques.

The key contributions of this work are as follows:

- We propose the first attention-based model for processing 4D spatio-temporal data. Based on the multi-head self-attention model Vaswani et al. (2017) and the pooling operation Devlin et al. (2018), we developed a self-attention model that attends to a series of embedded 3D frames.
- Our model achieves substantial improvement in classification performance over previous works on task state decoding using the 7-task benchmark from the HCP dataset.
- We demonstrate transfer learning of the learned features under different conditions and tasks.

2. Related Work

Recently, there has been much interest in applying machine learning to fMRI data Pereira et al. (2009). Convolutional neural networks (CNNs) LeCun et al. (1998) have been successfully adopted to learn representational features from fMRI data. Huang et al. (2017) proposed an architecture based on the sparse convolutional autoencoder to learn high-level features from handcrafted time series derived from the raw fMRI data. Wang et al. (2018) proposed a 4-layer CNN that classifies tasks from the raw fMRI voxel

values. Both works were demonstrated on HCP, treating the entire fMRI sequence as input.

More recent efforts explore the use of sequence models, that ingest each frame of the fMRI sequence in a sequential order. As such, sequence models are designed to capture temporal correlations between consecutive frames. [Mao et al. \(2019\)](#) used a CNN to extract features from each fMRI frame, then applied each frame’s feature as input into a Long Short-Term Memory (LSTM) [Hochreiter and Schmidhuber \(1997\)](#) model. Furthermore, combinations of LSTM and CNN have become a popular approach for fMRI analysis. [Thomas et al. \(2018\)](#) used a bidirectional LSTM [Schuster and Paliwal \(1997\)](#) in conjunction with a simple CNN. [Pominova et al. \(2018\)](#) used a LSTM in conjunction a residual network (ResNet) [He et al. \(2016\)](#).

3. Methods

Unlike RNNs and LSTMs that process data sequentially, transformers [Vaswani et al. \(2017\)](#) offer a different way to reason about sequence data, by *not* treating them as sequence, but rather, processing the entire sequence of data all at once. This allows frames that are temporally far apart to have the same level of access to each other’s information, in stark contrast to RNNs and LSTMs. This non-sequential access allows transformers to outperform RNNs in terms of long-term memory, as well as scalability due to its parallelizable architecture. Due to these advantages, we propose the first transformer-based model for 4D fMRI data, and demonstrate its utility in fMRI task state decoding.

3.1. Problem Statement

An instance of the fMRI data is a tuple $(\mathbf{x}_{1:T}, \mathbf{y})$, where $\mathbf{x}_{1:T} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ is a sequence of T 3-D frames. Each frame $\mathbf{x}_t \in \mathbb{R}^3$ is a snapshot of the subject’s brain activity at time t . The tasks and conditions

under which $\mathbf{x}_{1:T}$ was captured during the fMRI procedure are denoted as \mathbf{y} .

3.2. Frame Embedding

First, we need a model to transform each fMRI frame into a compressed numeric representation, or *embedding*.

3.2.1. 3D CNN

Using 3D CNN in fMRI task state decoding was proposed by [Wang et al. \(2018\)](#). The 3D kernels in the 4 convolutional layers are 1x1x1, 7x7x7, 3x3x3 and 3x3x3, respectively. Except for the first layer, whose stride is 1, all other layers adopt a stride of 2. This model uses global average pooling [He et al. \(2016\)](#) instead of max pooling on the final feature maps. The resulting feature maps are flattened (with dimension 512, which is the number of output channels in the CNN) and fed into a classifier with 2 fully-connected layers, whose output dimensions are 64 and 7, respectively. This model accounts for the temporal dimension of fMRI data by using 1x1x1 convolution (thus treating the temporal dimension as extra channels) to weigh every voxel accordingly across the time series. This model is implemented as a baseline method and will be referred to as CNN.

3.2.2. 3D RESNET

The 3D ResNet model [Pominova et al. \(2018\)](#) extends 3DCNN by adopting residual connections and bottleneck blocks from ResNet. Since no code was made available, we implemented the 3D ResNet model based on our interpretation of the model description given in [Pominova et al. \(2018\)](#) as follows: We augmented the 2D ResNet-18 [He et al. \(2016\)](#) into its 3D counterpart, the 3D ResNet-18 (cf. Table 1). This model is implemented as a baseline method and will be referred to as 3DResNet. For our proposed model, we make use of the spatial extractor, 3D ResNet-18, within the 3DResNet.

3.3. Temporal Modeling

3.3.1. LSTM-BASED MODELS

In Mao et al. (2019) and Pominova et al. (2018), an LSTM was used to model the temporal correlation between the embedded frames. The LSTM model has 2 hidden layers, each with 128 hidden units. In Thomas et al. (2018), a bidirectional LSTM was used instead. This 3DResNet-LSTM model – the combination of 3D ResNet and LSTM – is implemented as a baseline method.

3.3.2. SELF-ATTENTION AND TRANSFORMERS

We hypothesize that attention is valuable in modeling the temporal dynamics of fMRI data. The transformer Vaswani et al. (2017) applies self-attention to the entire sequence to model relationships between the frames in the sequence. These relationships are quantified as attention scores, which weight each frame due to its relevance to other frames. Through empirical experiments, we optimized the number of attention layers to 2 and attention heads to 8 for the transformer in our proposed model.

3.4. Our Transformer-based Model

Our model, called *Brain Attend and Decode* (BAnD) (cf. Figure 1), adapts the transformer model to the task of fMRI task state decoding. The transformer maps each embedded 3D frame $\mathbf{e}_t = \sigma(\mathbf{x}_t)$ into a latent vector \mathbf{z}_t that is conditioned with information from other embedded frames. To combine these latent vectors (also known as *transformed embeddings*) into a single representation that can be fed into the classifier, we adopt the approach in Devlin et al. (2018), whereby we insert a shared artificial token (\mathbf{e}_*) to the beginning of each series and train it together with the transformer. This can be understood as the one “frame” that ex-

Layer Name	3D ResNet-18 (kernel size, number of filters, stride)
conv1	$7 \times 7 \times 7, 64, \text{stride } 2$
conv2	$\begin{bmatrix} 3 \times 3 \times 3, 64, 1 \\ 3 \times 3 \times 3, 64, 1 \end{bmatrix} \times 2$
conv3	$\begin{bmatrix} 3 \times 3 \times 3, 128, 2 \\ 3 \times 3 \times 3, 128, 2 \end{bmatrix} \times 2$
conv4	$\begin{bmatrix} 3 \times 3 \times 3, 256, 2 \\ 3 \times 3 \times 3, 256, 2 \end{bmatrix} \times 2$
conv5	$\begin{bmatrix} 3 \times 3 \times 3, 512, 2 \\ 3 \times 3 \times 3, 512, 2 \end{bmatrix} \times 2$
global average pool	$1 \times 1 \times 1 \times 512$
flatten	embedding of size 512

Table 1: Base 3D ResNet-18 architecture

tracts essential information from the entire time series. By visualizing the attention of that frame on the other frames, this provides insights about which frames are important to the classification problem. The workflow of BAnD is summarized in Figure 1.

Similar to 3DResNet-LSTM, we apply the base 3D ResNet-18 to embed each frame of the data, i.e., $\mathbf{e}_t = \sigma(\mathbf{x}_t)$ for $t = 1, \dots, T$. A special pooling token \mathbf{e}_* is concatenated to this series, i.e., $\{\mathbf{e}_*, \mathbf{e}_1, \dots, \mathbf{e}_T\}$. Positional encoding Vaswani et al. (2017) is added to each embedding to enable position understanding. The transformer will map these embeddings $\{\mathbf{e}_*, \mathbf{e}_1, \dots, \mathbf{e}_T\}$ into latent vectors $\{\mathbf{z}_*, \mathbf{z}_1, \dots, \mathbf{z}_T\}$. The pooling token \mathbf{z}_* is fed into a classifier (comprised of 2 fully-connected layers with output dimensions of 512 and the number of classes) to get the logits for task classification. Finally, softmax is applied to get output probabilities for each class. We applied dropout with a rate of 0.2 to the transformer and classifier layers.

For BAnD and 3DResNet-LSTM, because we embedded each 3-D frame individually before passing the embeddings to the temporal model, limited GPU memory was an issue. By employing model parallelization, we were able to increase our batch size from 2 to 64, by training across 8 nodes utilizing a total of 32 GPUs during training.

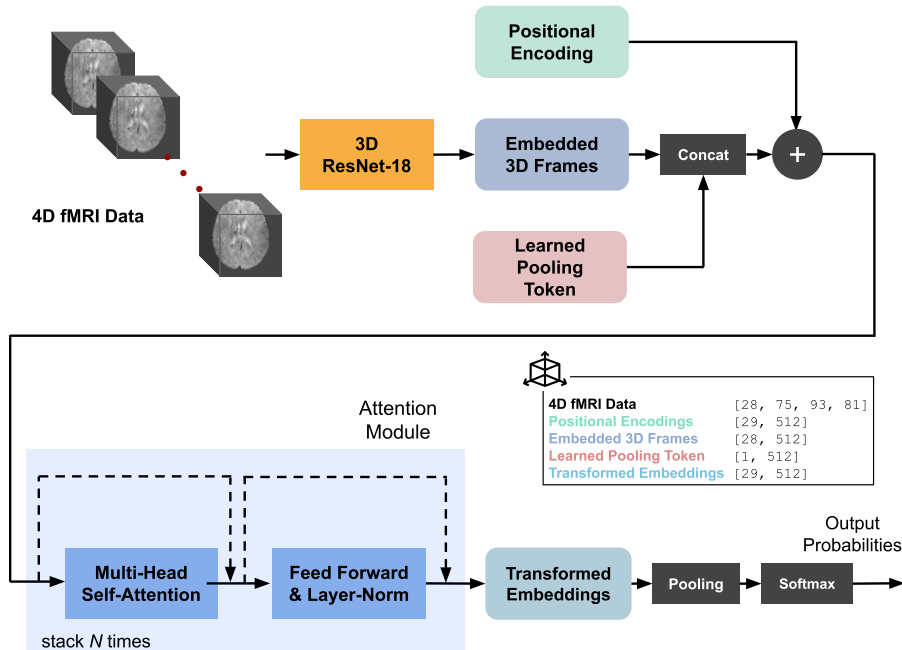


Figure 1: BANd architecture. Our model has $N = 2$. We show input dimensions at each component for the HCP 7-task dataset.

4. Datasets

4.1. HCP 7-task

The HCP-YA fMRI dataset includes 7 functional tasks, each under different conditions. These 7 tasks as a whole provide good brain activation coverage Barch et al. (2013), thus a classifier trained on this dataset would be useful for brain state decoding on a wide range of functional tasks. We used the same data as Wang et al. (2018) (cf. Table 2), which is a subset of the full HCP-YA fMRI data. Like Wang et al. (2018), we cropped out empty regions in the raw fMRI data, so the spatial dimensions $[d_x, d_y, d_z]$ are reduced from $[91, 109, 91]$ to $[75, 93, 81]$. The data from Wang et al. (2018) included 8 extra seconds post task/condition. We omitted these extra post frames (sampled at 0.72 frames per second) and examined only the frames during the actual task/condition. Depending on the task c , our data instance

varies in its dimensions $[T_c, d_x, d_y, d_z]$ where $[d_x, d_y, d_z]$ are same as above, but T_c ranges from 16 to 38.

Only 965 of the 1,034 subjects performed all 7 tasks in this benchmark. For each task, a subject might have performed more than one run. The total number of 4D data instances in this dataset is 17,368, with class distribution as shown in Table 2. The 17,368 instances were partitioned into train set (70%), validation set (10%) and test set (20%). Because each subject might perform a task multiple times, to prevent leakage of information across train/validation/test partitions, we stratified the splits so that data from a subject only belongs to one of the partitions.

4.2. Transfer learning

We distilled 3 subsets from HCP to evaluate transfer learning under intra-task (same task, different conditions) and inter-task (different tasks) conditions.

Task - Condition	# Frames	# Instances	Percentage
Emotion - Fear	25	2,895	16.7%
Gambling - Loss	38	1,930	11.1%
Language - Present story	28	3,860	22.2%
Motor - Right hand	16	1,930	11.1%
Relational - Relation	22	2,895	16.7%
Social - Mental	31	2,893	16.6%
WM - 2-back places	38	965	5.6%

Table 2: HCP 7-task dataset summary. Each instance is a complete 4D data point of a task/condition. A subject might perform a task multiple times, resulting in multiple instances. # Frames is the time dimension of each task/condition.

HCP 7-task-b set: This dataset was extracted from HCP and has similar characteristics as the HCP 7-task subset that we used to train BANd. This dataset is used for the intra-task transfer evaluation. The fMRI tasks are the same, but the conditions are different. Instead of the conditions in Table 2, we have: neutral, win, math, right foot, match, random, 2-back body, for the respective tasks. Pre-processing and augmentation scheme remain the same. There are 18,896 data instances in total.

Motor transfer set: In this setting, BANd was trained with data instances from the following 6 tasks: Working memory, Gambling, Relational, Social, Language, Emotion. Then the model was evaluated on the held-out Motor task, under 5 different conditions: right hand, left hand, right foot, left foot and tongue. There are 9,650 data instances in total.

WM transfer set: This dataset is similar to the Motor transfer set, except the Working Memory (WM) task is held out instead of Motor. There are 3,860 data instances in total.

4.3. Data Processing

It was important to normalize each data instance to have the same number of timesteps. Otherwise, the temporal model would simply learn to count frames to infer the task/condition. Following the data processing scheme in Wang et al. (2018), we ex-

tracted sets of k contiguous frames from each data instance (i.e., each set is $\mathcal{S} = \{\mathbf{x}_r, \mathbf{x}_{r+1}, \dots, \mathbf{x}_{r+k-1}\}$ where r is a random index) for training. During validation and testing, only the set corresponding to $r = 0$ (the first k frames) were selected. For data instances with less than k frames, we looped the time series until k frames are achieved, as is commonly done in video recognition Carreira and Zisserman (2017). In our experiments, $k = 28$.

5. Results and Analysis

In addition to the deep-learning-based baseline models described in section 3, we also compared against other traditional baselines that use parcellated fMRI data. Due to the lack of code and data processing scripts from previous works, we implemented all baseline models from scratch, with the intent to share our code to promote reproducibility within the research community. For linear models, we used Scikit-learn package¹. For neural network models, we used PyTorch² and adapted the transformer code from the PyTorch-Transformers repo³.

For all our experiments, we used an IBM Power8 computing cluster, where each node is equipped with 4 Nvidia P100 GPU cards.

1. <https://scikit-learn.org/>, version 0.22.1

2. <https://pytorch.org/>, version 1.2

3. <https://github.com/huggingface/pytorch-transformers>

5.1. HCP 7-task

5.1.1. COMPARISON WITH METHODS USING PARCELLATED FMRI DATA

We validated our model against a regression-based baseline that utilizes parcellated fMRI data. We used the parcellated data that came with the HCP-YA dataset, which had undergone the Glasser parcellation [Glasser et al. \(2016a\)](#). Each data instance contains 360 cortical and 14 subcortical regions. Prior to computing the correlation matrix, we applied global signal regression to subtract out common trends across the timeseries. This is performed using a 2 parameter linear regression capturing shift and scale perturbations from the global average. We then constructed a connectome matrix from the fMRI data [Behrens and Sporns \(2012\)](#). We next utilized the Pearson Correlation matrix between parcellated regions to represent the amount of functional connectivity between regions [Sporns \(2011\)](#). Then, we band-filtered the timeseries signals to the range 0.009 Hz - 0.25 Hz using a 4-th order Butterworth filter. The correlation step results in a 374×374 matrix with values ranging from -1 to 1.

We performed this regression task in two ways: logistic regression and fully connected neural network. For logistic regression, we experiment with non-regularized model as well as an L1 regularized version with k-fold cross-validation. We set k to be 3 and tuned the inverse of regularization strength in the range of 10^c with c in $[-3, -2, \dots, 3]$. As for the neural network model, we build a 3-layer fully-connected network with ReLU activation. Its output is fed into a softmax layer and the model is trained with cross-entropy loss. We did light hyper-parameter tuning and the size of the hidden layer was set to 320.

5.1.2. COMPARISON WITH METHODS USING NON-PARCELLATED FMRI DATA

Results for each task are reported with means and standard deviations from 3 different random splits of training, validation and testing data, whenever applicable. For neural network models, we used Adam optimizer [Kingma and Ba \(2014\)](#) with weight decay of 0.0001. We empirically finetuned the learning rates for each model and used a cosine annealing learning rate schedule.

BAnD outperforms all other models, as shown by the increasing accuracy values in [Figure 3](#): from 91.4% with 3DCNN, to 94.5% with 3DResNet, to 95.1% with BAnD. We conjecture that this increase in performance stems from the added attention capability to understand complex spatio-temporal dynamics in 4D fMRI data. It is also important to note that with a logistic regression model, correlation matrix regression approach with L1 regularization achieved an accuracy of 92.2%, which is better than 91.4% of the 3DCNN model. However, extensive pre-processing is required to parcellate raw fMRI data and compute correlation matrices. Performance at down-stream tasks might also depend on different parcellation strategies. It is further unclear how to extract/analyze the temporal dynamics inherent in fMRI data with this approach.

As for the 3DResNet-LSTM model, even though we experimented with an extensive hyperparameter sweep of the number of hidden layers $[1, 2, 3]$, the number of hidden units for each layer $[64, 128, 256]$, the directionality (bi- or uni-) of the LSTM model, as well as different dropout rates $[0.2, 0.4, 0.6]$, the model did not converge to a satisfactory validation/test accuracy. With limited computational budget and time, we trained each of these configurations for 24 hours under a distributed environment with 32 GPU cards.

Further work is needed to analyze why these LSTM models did not converge under this HCP benchmark. Nonetheless, our BAnD model was able to converge in under 12 hours with the same computational setting. This might be due to the Transformer’s advantage of parallelizability and faster convergence, as suggested by Vaswani et al. (2017).

Furthermore, as Popel and Bojar (2018) noted, a bigger batch size (up to a certain threshold) yields better results with the Transformer model. Due to limited GPU memory, BAnD was trained with a batch size of 2 on each GPU card. To increase the batch size, we froze the 3D ResNet-18 layers in the converged BAnD model and trained a new Transformer on top. With this trick, we were able to increase the batch size from 2 to 128, thus deriving at a new model called BAnD++. From Table 3, BAnD++ shows a significant increase in classification accuracy, from 95.1% to 97.2%.

To evaluate the results qualitatively, we applied t-SNE Maaten and Hinton (2008) to visualize the final embeddings generated by the 3D ResNet and BAnD++ models. As can be seen in Figure 2, 3DResNet, though quite effective, was not able to disambiguate clearly the data instances belonging to Emotion, Gambling and Social tasks. On the other hand, BAnD++, with the added capability of a Transformer, was able to separate out those instances into separate groups.

Interestingly, even though the joint model 3DResNet-LSTM did not converge to satisfactory classification performance in our experiments, a model with a 3D ResNet-18 spatial model pre-trained as BAnD, with an LSTM temporal head (denote this model as 3DResNet-LSTM++) achieves similar results to BAnD++. This shows that BAnD with the Transformer head instead of an LSTM head not only made it possible for the spatial feature extractor to converge faster but the learned features are also general

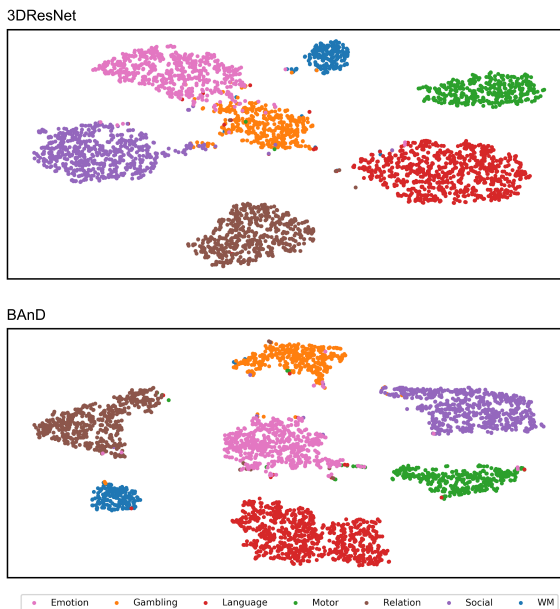


Figure 2: 3DResNet and BAnD embeddings t-SNE. For 3DResNet, which does not model temporality explicitly, the clusters are more tangled, while for BAnD, which learns the temporal dynamics in 4D fMRI through Transformer, the clusters are more separated, resulting in better classification performance.

enough for a new LSTM to learn the temporal patterns of the 4D fMRI data series, and achieves significant performance.

5.2. Transfer learning

In this section, we discuss the experimental results on the transferability of BAnD features under the intra-task and inter-task settings, as described above.

Intra-task, HCP 7-task-b set: The goal of this subset is to investigate whether the features learned with BAnD can transfer to a similar dataset. We first froze the spatial feature extractor (i.e., the 3D ResNet-18 layers) in BAnD trained on HCP 7-task dataset, added a new temporal Transformer head (with the same hyperparameters) and trained it on HCP 7-task-b dataset, without

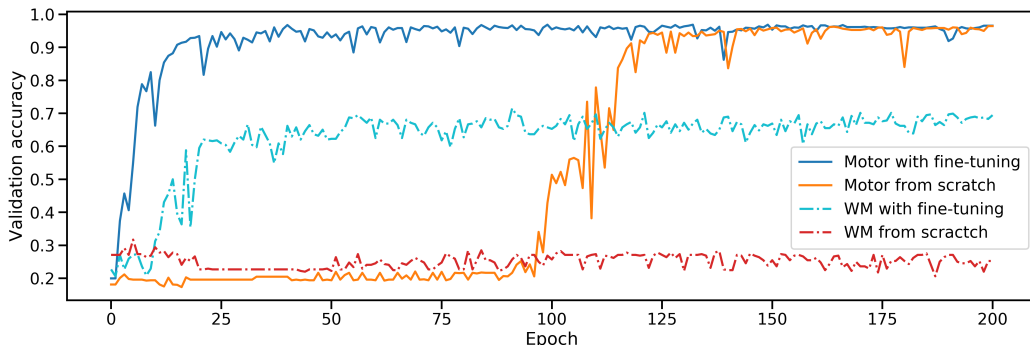


Figure 3: Finetuning vs training from scratch

Logistic regression	0.915 ± 0.003
Logistic regression (L1 reg)	0.922 ± 0.002
Fully-connected NN	0.915 ± 0.004
3DCNN	0.914 ± 0.0204
3DResNet	0.948 ± 0.011
3DResNet-LSTM	0.534 ± 0.053
BAnD (ours)	0.951 ± 0.0062
BAnD++ (ours)	0.972 ± 0.0057
3DResNet-LSTM++	0.970 ± 0.0037

Table 3: Accuracy results for different models on the HCP 7-task. Means and standard deviations are calculated from 3 different splits of training, validation and testing data.

finetuning the 3D ResNet-18 layers. As can be seen in Table 4, our model achieved 93.6% accuracy in this setting, which suggests that the spatial features learned with BAnD are highly applicable to similar tasks.

Inter-task, Motor transfer: Under this setting, the test tasks are significantly different from the transfer tasks. With the same approach as applied to the HCP 7-task-b, where we froze the 3D ResNet-18 layers in BAnD and trained a new Transformer head, we observed an accuracy of 51.1%. This result suggested that there was room for improvement. So we further finetuned the 3D ResNet-18 layers to the task of classifying the different Motor conditions. As shown in 4,

7-task-b	0.936 ± 0.012
Motor tasks	0.511 ± 0.027
WM tasks	0.623 ± 0.031
Motor tasks with finetuning	0.941 ± 0.0014
Motor tasks from scratch	0.934 ± 0.0067
WM tasks with finetuning	0.71 ± 0.0043
WM tasks from scratch	0.23 ± 0.018

Table 4: Transfer learning and finetuning accuracy results

we observed a significant increase in classification performance, from 51.1% to 94.1%. On the other hand, a model trained from scratch achieved 93.4% classification accuracy, but took significantly more time, while the *Motor with finetuning* model was able to leverage features learned from the tasks it was pre-trained on.

Inter-task, Working Memory transfer: We repeat the experiment above with the held-out WM tasks instead of Motor tasks. We noticed the same patterns as observed in the unseen Motor tasks, in which further finetuning helped our model achieve a significant increase in performance and with shorter training time.

5.3. Analysis

Temporal activation maps: Attention models such as Transformer allows one to examine the attention weights of the model to

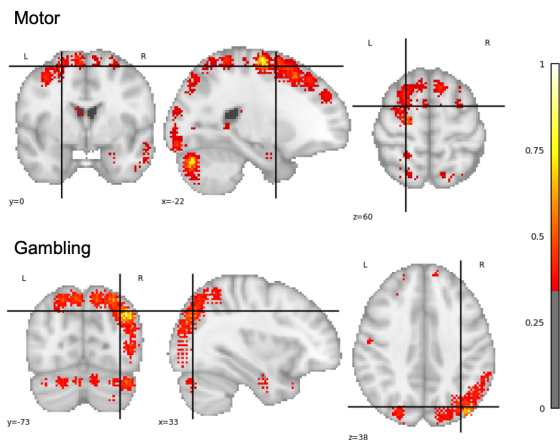


Figure 4: BAnD spatial activation map for Motor (Right hand) and Gambling (Loss) task

see which frame is weighted more for a classification result. However, our base BAnD model was trained with 8 attention heads, which makes it hard to discern which set of attention weights is truly important for a classification. Thus, we trained a new BAnD++ with only 1 attention head to make it easier to visualize. To show attention across more frames, we visualize the model’s attention weights across $k = 32$ frames. We calculated the attention weights across time for each data instance from the test set that was correctly classified by the model. We then averaged all attention weights in a task to get a representative set of weights. Figure 5 shows the result for the 7 tasks in HCP 7-task benchmark. Interestingly, across different tasks, the 7th frame of a data series seems to have a strong influence on the predictions. This can be due to a delay in the BOLD signal in response to a stimuli [Liao et al. \(2002\)](#).

Spatial activation maps: Using Grad-Cam [Selvaraju et al. \(2017\)](#), we were able to find the set of voxels that were highly activated for a certain class prediction by our BAnD model. We calculated this spatial activation map for each data series, then averaged across data points for a class and

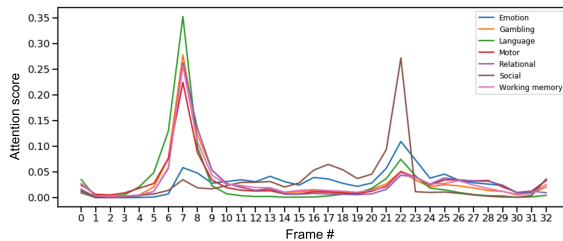


Figure 5: BAnD temporal activation map

picked the 7th frame in each series because our temporal analysis showed that 7th frame was deemed important by the model. We used the `nilearn`⁴ package to project from 3D voxel space to brain statistical maps. Figure 4 shows those for the Motor and Gambling task.

To motivate reproducible research in the area of fMRI task state decoding, we plan to release our data processing scripts, codes and pretrained models upon acceptance.

6. Conclusion

In this work, we presented a novel attention-based model for processing 4D fMRI data, and showed that our proposed novel architecture, BAnD, which is based on a multi-headed self-attention module, achieves significant performance gains compared to previous works on a recent 7-task fMRI benchmark from the Human Connectome Project dataset. We further demonstrated transfer learning capability of BAnD’s learned features to unseen conditions and tasks and show that BAnD achieves competitive results with finetuning. To try to understand what BAnD was attending to for its classifications, we computed both spatial and temporal activation maps, highlighting brain regions and frame important for task decoding. Future work includes transferring BAnD to other fMRI data sets, such as ADHD, or to other modalities of medical imaging.

4. <https://nilearn.github.io/>

Acknowledgments

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and was supported by the LLNL LDRD Program under Project No. 19-ERD-009. LLNL-CONF-788817.

References

- Deanna M Barch, Gregory C Burgess, Michael P Harms, Steven E Petersen, Bradley L Schlaggar, Maurizio Corbetta, Matthew F Glasser, Sandra Curtiss, Sachin Dixit, Cindy Feldt, Dan Nolan, Edward Bryant, Tucker Hartley, Owen Footer, James M Bjork, Russ Poldrack, Steve Smith, Heidi Johansen-Berg, Abraham Z Snyder, David C Van Essen, and WU-Minn HCP Consortium. Function in the human connectome: task-fMRI and individual differences in behavior. *Neuroimage*, 80:169–189, October 2013.
- Timothy E J Behrens and Olaf Sporns. Human connectomics. *Curr. Opin. Neurobiol.*, 22(1):144–153, February 2012.
- Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Matthew F Glasser, Timothy S Coalson, Emma C Robinson, Carl D Hacker, John Harwell, Essa Yacoub, Kamil Ugurbil, Jesper Andersson, Christian F Beckmann, Mark Jenkinson, Stephen M Smith, and David C Van Essen. A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615):171–178, August 2016a.
- Matthew F Glasser, Stephen M Smith, Daniel S Marcus, Jesper L R Andersson, Edward J Auerbach, Timothy E J Behrens, Timothy S Coalson, Michael P Harms, Mark Jenkinson, Steen Moeller, Emma C Robinson, Stamatios N Sotiropoulos, Junqian Xu, Essa Yacoub, Kamil Ugurbil, and David C Van Essen. The human connectome project’s neuroimaging approach. *Nat. Neurosci.*, 19(9):1175–1187, August 2016b.
- Michael D Greicius, Gaurav Srivastava, Allan L Reiss, and Vinod Menon. Default-mode network activity distinguishes alzheimer’s disease from healthy aging: evidence from functional mri. *Proceedings of the National Academy of Sciences*, 101(13):4637–4642, 2004.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Heng Huang, Xintao Hu, Yu Zhao, Milad Makkie, Qinglin Dong, Shijie Zhao, Lei Guo, and Tianming Liu. Modeling task fmri data via deep convolutional autoencoder. *IEEE transactions on medical imaging*, 37(7):1551–1561, 2017.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-

- based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Chien Heng Liao, Keith J Worsley, J-B Poline, John AD Aston, Gary H Duncan, and Alan C Evans. Estimating the delay of the fmri response. *NeuroImage*, 16(3):593–606, 2002.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- Zhenyu Mao, Yi Su, Guangquan Xu, Xueping Wang, Yu Huang, Weihua Yue, Li Sun, and Naixue Xiong. Spatio-temporal deep learning method for adhd fmri classification. *Information Sciences*, 499:1–11, 2019.
- Thomas Naselaris, Kendrick N Kay, Shinji Nishimoto, and Jack L Gallant. Encoding and decoding in fmri. *Neuroimage*, 56(2):400–410, 2011.
- Francisco Pereira, Tom Mitchell, and Matthew Botvinick. Machine learning classifiers and fmri: a tutorial overview. *Neuroimage*, 45(1):S199–S209, 2009.
- Marina Pominova, Alexander Bernstein, Alexey Artemov, Evgeny Burnaev, Maxim Sharaev, and Ekaterina Kondratyeva. Voxelwise 3d convolutional and recurrent neural networks for epilepsy and depression diagnostics from structural and functional mri data. 09 2018. doi: 10.1109/ICDMW.2018.00050.
- Martin Popel and Ondřej Bojar. Training tips for the transformer model. *The Prague Bulletin of Mathematical Linguistics*, 110(1):43–70, 2018.
- Serge ARB Rombouts, Jessica S Damoiseaux, Rutger Goekoop, Frederik Barkhof, Philip Scheltens, Stephen M Smith, and Christian F Beckmann. Model-free group analysis shows altered bold fmri networks in dementia. *Human brain mapping*, 30(1):256–266, 2009.
- U Sabatini, K Boulanouar, N Fabre, F Martin, C Carel, C Colonnese, L Bozzao, I Berry, JL Montastruc, F Chollet, et al. Cortical motor reorganization in akinetic patients with parkinson’s disease: a functional mri study. *Brain*, 123(2):394–403, 2000.
- M. Schuster and K.K. Paliwal. Bidirectional recurrent neural networks. *Trans. Sig. Proc.*, 45(11):2673–2681, November 1997. ISSN 1053-587X. doi: 10.1109/78.650093. URL <https://doi.org/10.1109/78.650093>.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 618–626, 2017.
- Olaf Sporns. The human connectome: a complex network. *Ann. N. Y. Acad. Sci.*, 1224:109–125, April 2011.
- Yukiko Tada. Motor association cortex activity in parkinson’s disease—a functional mri study. *Rinsho shinkeigaku= Clinical neurology*, 38(8):729–735, 1998.
- Armin W. Thomas, Hauke R. Heekeren, Klaus-Robert Müller, and Wojciech Samek. Interpretable lstms for whole-brain neuroimaging analyses. *CoRR*, abs/1810.09945, 2018. URL <http://arxiv.org/abs/1810.09945>.
- David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy E J Behrens,

- Essa Yacoub, Kamil Ugurbil, and WU-Minn HCP Consortium. The WU-Minn human connectome project: an overview. *Neuroimage*, 80:62–79, October 2013.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- Xiaoxiao Wang, Xiao Liang, Yawen Zhou, Yanming Wang, Jin Cui, Huijuan Wang, Yu Li, Benedictor Alexander Nguchu, and Bensheng Qiu. Task state decoding and mapping of individual four-dimensional fmri time series using deep neural network. *arXiv preprint arXiv:1801.09858*, 2018.
- Andrew Webb and George C Kagadis. Introduction to biomedical imaging. *Med. Phys.*, 30(8):2267–2267, August 2003.
- Dong Wen, Zhenhao Wei, Yanhong Zhou, Guolin Li, Xu Zhang, and Wei Han. Deep learning methods to process fmri data and their application in the diagnosis of cognitive impairment: a brief overview and our opinion. *Frontiers in neuroinformatics*, 12:23, 2018.