Convolutional Networks for EEG Signal Classification in Non-Invasive Brain-Computer Interfaces

Colorado State University
Brain-Computer Interfaces Lab.
CEBL3

Elliott Forney, Charles Anderson, William Gavin and Patricia Davies

NSF
Grant Number 1065513.

Non-Invasive Brain-Computer Interfaces

Brain-Computer Interfaces (BCI) are systems for establishing a direct channel of communication between the human brain and a computerized device.

BCI utilize a communication protocol along with signal processing and machine learning algorithms to convey the user's intent to a computer system.

In our lab, we use Electroencephalography (EEG) to monitor brain activity because it is non-invasive, portable and relatively affordable.

An important application for BCI is in the development of assistive technologies for people with severe motor impairments.

For those who find it difficult to communicate or perform everyday tasks, even a somewhat slow BCI may prove to be an invaluable tool.

Synchronous Paradigms

Current BCI typically follow one of two paradigms.

Synchronous paradigms rely on the brain's response to a time-locked stimulus, known as an Event-Related Potential (ERP).

For example, the P300 speller type of BCI flashes the rows and columns of a virtual keyboard. The BCI then attempts to identify the letter that the user is attending to by analyzing the ERPs in the user's EEG signals following each stimulus presentation.

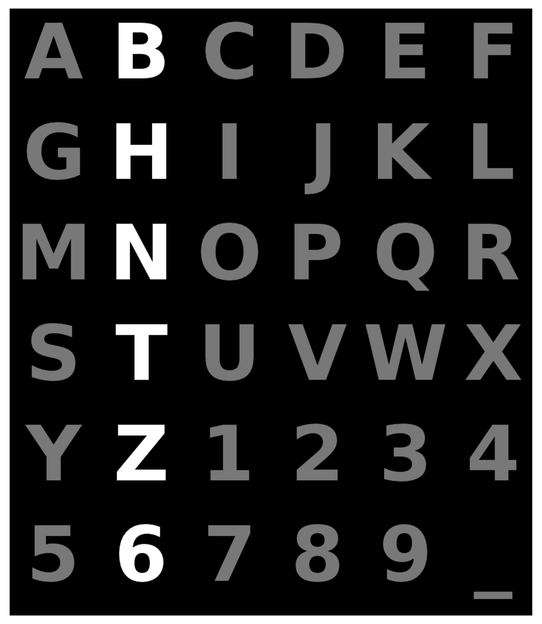

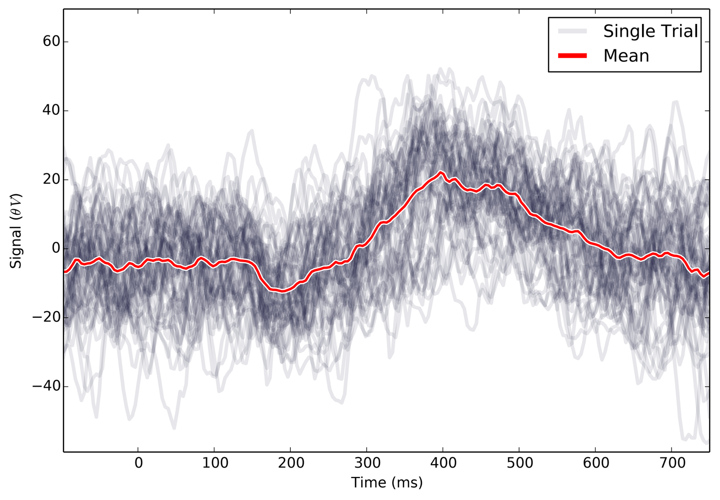

In order to detect the desired ERP, a number of trials are typically segmented by the stimulus onset and then averaged together.

Averaging smooths noise, artifacts, and trial-to-trial variations.

However, averaging may also remove useful information.

Asynchronous Paradigms

Asynchronous BCI paradigms are self-paced and typically do not require external stimuli.

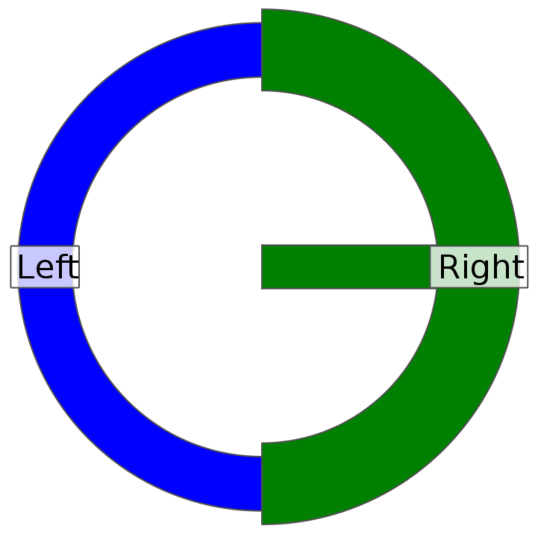

For example, in a motor imagery paradigm, a user might imagine moving their left arm to move a mouse cursor left or imagine moving their right arm to move the cursor to the right.

Amplitude changes in the sensory-motor regions of the brain can then be used to identify which imagined movement the user is performing.

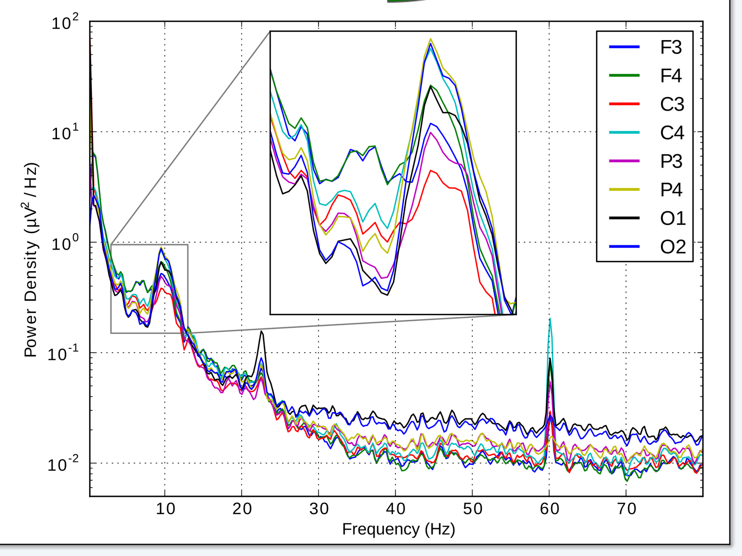

Since these changes are not time-locked, an appropriate model must be shift-invariant. Typically, this is handled by viewing the power spectrum of the signal.

These frequency changes are known as Event-Related (De)Synchronization (ERD/S).

Deep Learning

The current trend in machine learning is to move away from hand-crafted models and feature selection procedures in favor of algorithms that are capable of learning hierarchical representations.

These "deep networks" rely on minimal, if any, prior assumptions about the structure of the patterns in the data. In this case, the EEG signals.

Avoiding prior assumptions about the data may allow a machine learning algorithm to exploit patterns that are not currently known or well-understood by human researchers

Analysis of deep networks may also lead to new insights into the types of patterns that are useful for discriminating between different EEG signals.

Convolutional Networks

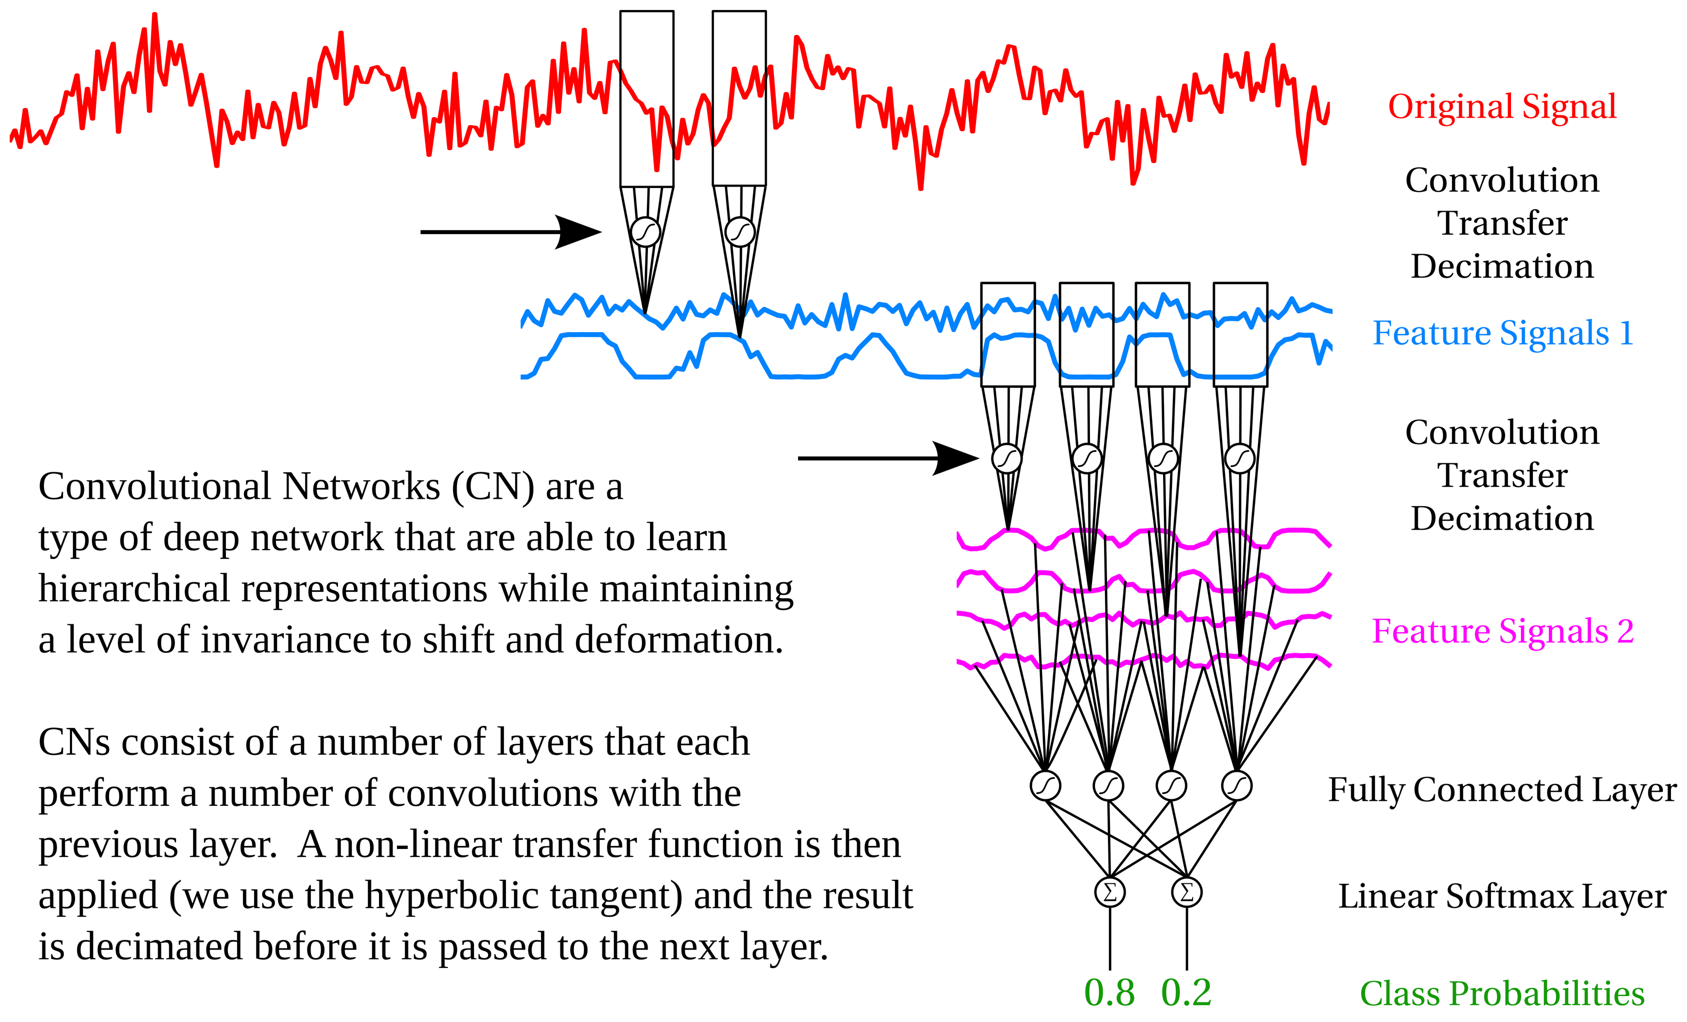

Convolutional Networks (CN) are a type of deep network that are able to learn hierarchical representations while maintaining a level of invariance to shift and deformation.

CNs consist of a number of layers that each perform a number of convolutions with the previous layer. A non-linear transfer function is then applied (we use the hyperbolic tangent) and the result is decimated before it is passed to the next layer.

The output of the convolutional layers is then passed to a fully connected network. This network contains a single non-linear layer followed by a linear layer with softmax readouts that indicate the probability that the signal belongs to each class.

The weights at each layer are learned using backpropigation, we use Scaled Conjugate Gradients.

Since the weights in each convolutional layer essentially "scan" the decimated output of the previous layer, these weights are considered to be shared.

This results in fewer parameters to be optimized than would be found in a similar network that is fully-connected.

This also allows the network to identify features, or events, that occur in the signal, regardless of the time at which they occurred, i.e., the network can learn shift-invariant features.

CN are becoming increasingly popular and have been shown to be very effective for some problems.

Receptive Fields and FIR Filters

The weights of the kernel of each convolution are often interpreted as the receptive field of an artificial neuron, a concept similar to structure often found in biological neural networks.

From this viewpoint, CN are designed to form a hierarchy of highly-specialized neurons with potentially overlapping receptive fields.

An alternate interpretation of the one-dimensional convolutional network is a series of non-linear Finite Impulse-Response (FIR) decimation filters.

From this viewpoint, each convolutional layer can be thought of as a filter-bank that processes the signal or selects features that are useful for classification.

Regularization

In order to prevent our CNs from fitting noise or unrelated patterns in our relatively small EEG datasets, we seek to limit the complexity of our models, i.e., our CNs are regularized.

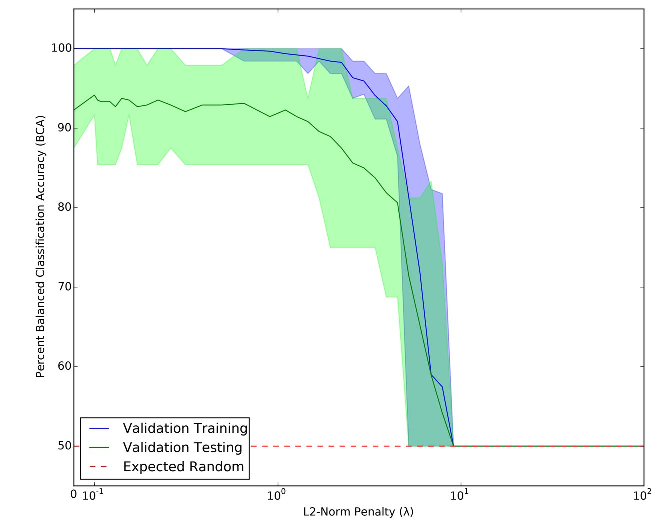

Currently, this is performed using an L2-norm weight decay with the same parameter at each layer.

Although this approach is effective, it is somewhat difficult to achieve models with moderate complexity. Further investigation into other regularization approaches appears to be necessary.

Preliminary Results

Subject	LDA BCA	LDA (AUC)	KNN BCA	KNN (AUC)	FN BCA	FN (AUC)	CN BCA	CN (AUC)
08	70.73	(61.46)	68.23	(65.62)	65.54	(56.88)	67.92	(64.79)
10	77.67	(68.33)	71.42	(52.50)	80.00	(63.33)	71.83	(62.50)
11	89.58	(79.17)	91.67	(78.33)	88.67	(75.00)	93.33	(81.67)
12	64.50	(57.50)	45.00	(48.33)	64.08	(53.33)	58.83	(61.67)
13	73.17	(66.67)	60.25	(60.00)	76.50	(65.00)	84.25	(77.50)
15	80.75	(73.33)	64.42	(59.17)	75.50	(65.00)	88.33	(77.50)
16	90.92	(81.67)	55.17	(53.33)	91.17	(77.50)	80.83	(72.50)
20	80.67	(72.50)	64.92	(52.50)	77.50	(65.83)	86.25	(75.83)
21	85.17	(80.83)	78.50	(69.17)	87.50	(75.83)	86.50	(80.83)
22	89.42	(85.83)	86.67	(73.33)	86.00	(87.50)	90.75	(80.00)
23	78.25	(65.83)	78.42	(74.17)	83.50	(70.83)	91.42	(73.33)
24	92.50	(74.17)	78.83	(66.67)	91.17	(77.50)	88.50	(77.50)
25	71.58	(65.00)	68.00	(55.83)	73.42	(73.33)	76.25	(67.50)
26	83.83	(80.00)	60.58	(61.67)	80.33	(76.67)	88.17	(77.50)
27	93.58	(90.83)	87.58	(66.67)	94.33	(85.83)	95.33	(84.17)
28	71.08	(60.83)	65.58	(65.00)	71.92	(65.00)	70.08	(66.67)
Mean	80.84	(72.75)	70.33	(62.64)	80.32	(70.90)	82.41	(73.84)

We have recently run several preliminary experiments that have yielded mixed but encouraging results.

In offline experiments using a synchronous P300-style task, CNs outperform regularized Linear Discriminant Analysis, K-Nearest Neighbors, and a traditional Feed-forward Network by about 3% Balanced Classification Accuracy or 2% Area Under the ROC Curve (AUC). This is a small but notable improvement for initial experiments.

For asynchronous Mental-Task BCI, our training performance is currently near 100% test performance near 50% for four tasks. This suggests that overfitting is a significant challenge that must be overcome in order for this type of BCI to work effectively with the small amount of data that can reasonably be collected during an interactive calibration phase.

Future Work

Improved methods for hyper-parameter selection (number and size of layers) and regularization appear to be the most important next step:

A random or meta search may allow for the weight decay to be adjusted at each layer. Introducing noise into the input layer or using "dropout" techniques are alternative regularization techniques that should be considered.

Variations in the networks architecture are also interesting avenues to explore. For instance, replacing the convolution operations with recurrent networks may yield models with fewer parameters and better generalization.

Of course, the final goal of this work is to produce usable BCI that perform well. Future works should focus on producing CNs that are fast enough to be used in interactive BCI and that are robust to the challenges found in real-world use cases.