Mapping of Hierarchical Activation in the Visual Cortex

Suman Chakravartula, Denise Jones, Guillaume Leseur

CS229 Final Project Report. Autumn 2008.

Introduction

There is much that is unknown regarding the communication between different areas in the brain. This is a research area of growing interest with the realization of applying technological advances to this field. In this paper we will describe a method to apply machine learning to map communication between regions of the Visual Cortex (VC).

Neural activity in the VC is hierarchical, and activation due to stimuli at one area leads to corresponding activity in specified regions of another area [1]. This can be illustrated by viewing the activation in voxels of each region over a given period of time using fMRI data. It has been shown that the signals for vision pass through the first area, V1, before entering other areas such as V2, and there is a mirrored image of the VC on each side of the brain, which we will refer to as RV1, RV2 and LV1, LV2 (we will treat these as separate regions of interest, ROIs, in our mapping). In addition, we ignore any possible cross-wiring between hemispheres in these regions. Given the hierarchical processing of information from the stimulus to V1 and then to V2, our intent is to build a model to confidently predict mappings between these regions. In the process, we also build a model to predict V1 from the stimulus.

Data and Methods

The FMRI data for our project was provided by Professor Brian Wandell's lab at Stanford University. It is preprocessed to remove noise from the scanner and compensate for minor head movements of the subject. Our training sets and test sets consist of 1300 and 200 data points respectively. For example, in the model of stimulus to V1 prediction, our training set consists of the stimulus over 1300 time points and corresponding activation of a single voxel in V1. Given the stimulus at a specified time point, the response variables are the voxel activation numbers at that time point.

We analyze the data by implementing a special type of linear regression called LASSO Regression. In addition, we implement regression using SVM and compare it with results from LASSO. We also compare our stimulus to V1 model with the pre-existing model of professor Wandell's lab. We then build a model for V1 to V2 prediction. Since the data is very high-dimensional, we implemented Principle Component Analysis (PCA) to reduce computational costs. By using PCA, we were also able to interpret the parameters coherently and visualize the mappings more clearly.

LASSO Regression

LASSO, or "least absolute shrinkage and selection operator," is a technique which improves upon ordinary least squares (OLS) regression by minimizing

$$\sum_{i=1}^{N}(y^{(i)} - \vartheta^T x^{(i)})^2 + \lambda \| \vartheta \|_1$$

where λ is the tuning parameter. It is implemented in a way as to shrink some $\vartheta^{(i)}$s and set others to 0. For example, if λ' is the normalization in the case of OLS and we set $\lambda = \lambda'/2$, then half the coefficients shrink to zero [2]. We choose to implement LASSO for the ease of interpreting the parameter ϑ to construct the mapping. We use 10-fold cross validation in order to determine the best λ and then retrieve ϑ from that model. This is done using lars package of R [3].

Once we have ϑ for the prediction of a given LV1 voxel from the stimulus, we reconstruct the stimulus by setting the value of each pixel *i* in the stimulus equal to ϑ^i. We then threshold the pixel values and set all pixels below the median value to the minimum value. By doing so, we obtain a receptive field mapping for a specific LV1 voxel as shown in Fig. 1(a). The corresponding receptive field given by the pre-existing model is shown in Fig. 1(c). By using LASSO regression, we are able to conveniently interpret ϑ to map the receptive field of any given voxel. We use the same procedure for the V1 to V2 model. Fig. 3(c) shows the strongly correlating voxel group (bright yellow) for an arbitrarily selected LV2 voxel. By this result, we can hypothesize that this group of LV1 voxels drive the selected LV2 voxel.

Fig. 2(a) and 3(a) show the actual and predicted signals together for the test data, which shows that the model predicts the voxel activation signal within a small margin of error. Generalization errors listed in Table 1 also show that the model is reasonably accurate. Therefore, in addition to providing interpretable results, the model provides a good prediction.

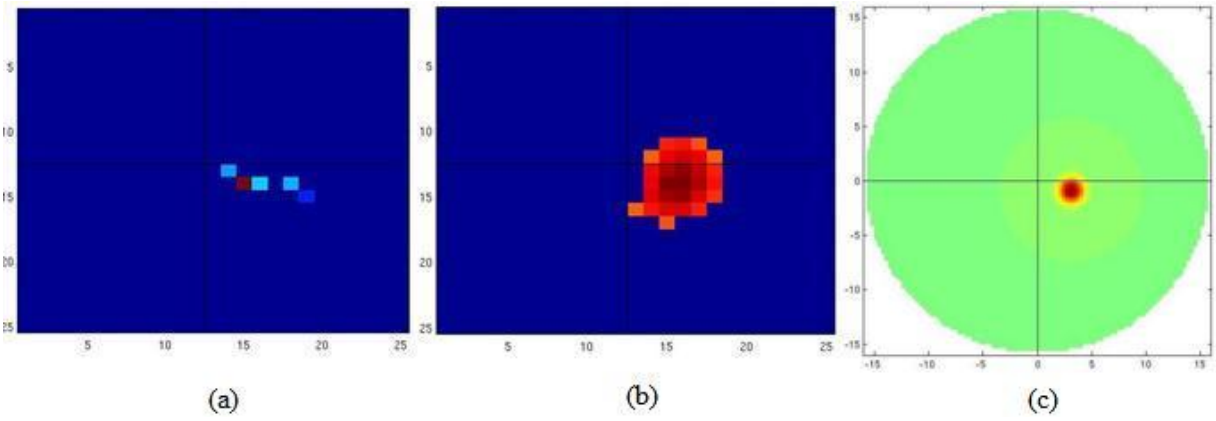

Fig.1. Receptive field mapping of a LV1 voxel according to parameters of (a) LASSO regression before PCA. (b) LASSO regression after PCA (c) pre-existing model of professor Wandell's lab.

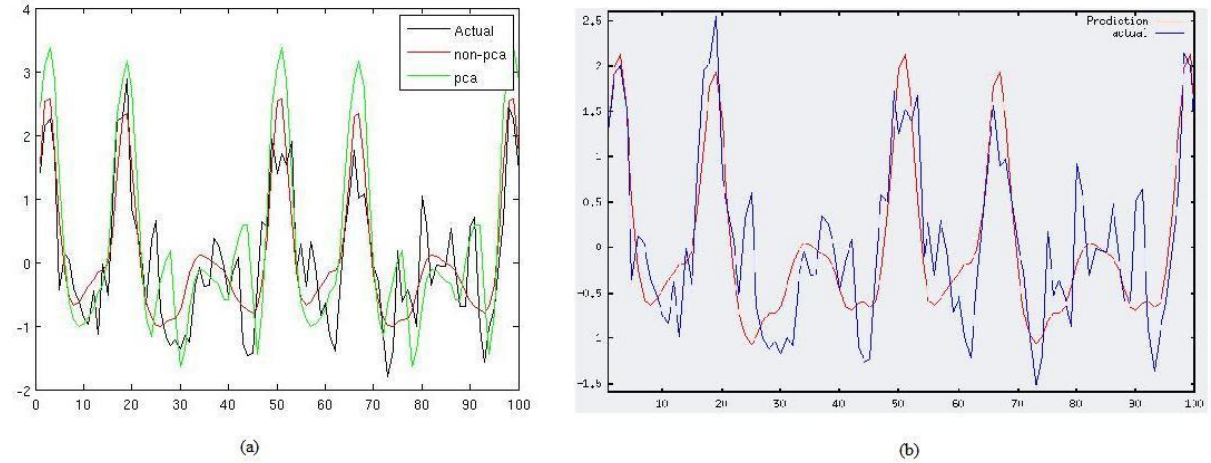

Fig. 2. Prediction curve plotted together with the actual signal for a LV1 voxel using the first 100 time points of the test data set, according to (a) LASSO regression (b) SVM regression.

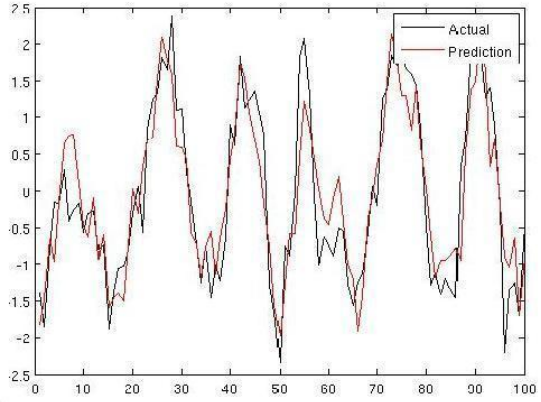

(a)

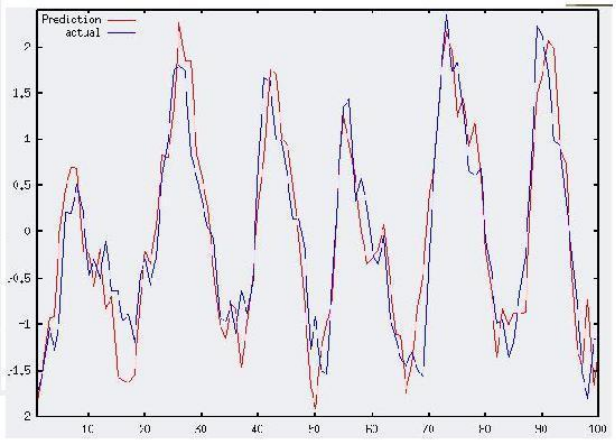

(b)

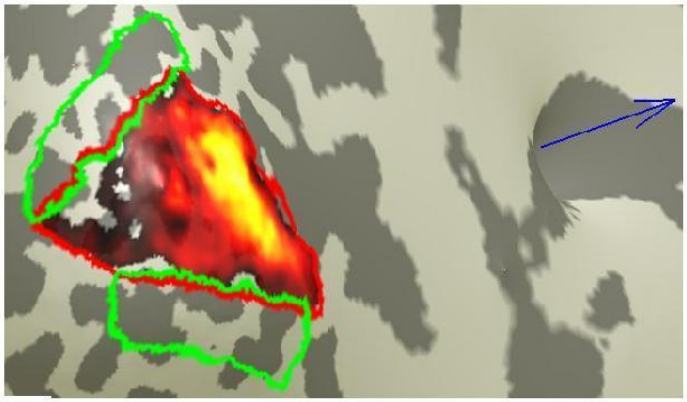

(c)

Fig. 3. V1 to V2 prediction curve plotted together with the actual signal for a LV1 voxel using the first 100 time points of the test data according to (a) LASSO regression and (b) SVM regression.

Fig. 3. (c) Mapping of V1 voxels(in bright yellow) for an arbitrarily selected V2 voxel (not shown but somewhere withing green boundaries) according to the LASSO regression model. The image has been cropped to show a small section of the inflated left hemisphere(as indicated by the blue arrow)

Model	Stimulus -> LV1	Stimulus -> RV1	LV1 -> LV2
LASSO	.5589	.3671	NA
LASSO after PCA	1.1064	.7054	.2462
SVM	.4587	.3105	NA
SVM after PCA	.4704	.3220	.17

Table 1. Generalization errors obtained by averaging 10 distinct voxel regressions per model.

SVM Regression

As another method to predict activation from the stimulus to V1 or from V1 to V2 we also applied a SVM regression algorithm. We decided to do so in order to have another prediction to compare with the LASSO regression results. We chose SVM for its capacity to handle high dimensional data and its renowned efficiency. We used the SVM_Light software that can be used for both classification and regression.

PCA

The facts that our data are highly dimensional, thousands of dimensions per voxel, and that these dimensions are highly correlated (nearby voxels have a very similar activity and the stimuli are simple geometric shapes) strongly encourage the use of PCA before applying the regression algorithms. PCA proved to be useful in two ways, namely, dimensionality reduction and coherence of mapping.

Contrary to the stimulus, it is not possible to downsample V1 data without losing much information. The computational cost of predicting the activation in a whole area from the whole stimulus is so high, that reducing the input dimension is a necessity. Upon performing PCA and plotting the decrease in component variance, we find that it drops very close to zero after about 600 PCs. With this 90% reduction in dimensionality, the computational cost of each regression lowers drastically. Previous figures show that we still achieve very efficient prediction within a reasonable period of time.

Fig. 1(a) and Table 1 show that there is some loss in predictability after PCA for stimulus to V1 prediction, particularly for the LASSO regression. However, the receptive field mapping is substantially more coherent and accurate compared with existing models as shown in Fig. 1(b). This is because these nearby voxels have a similar behavior and therefore they have a close decomposition in the new basis (the principal component basis). This improves the coherence of the mapping and its visualization.

Results

We have results for the two types of predictions and mapping we worked with: stimulus to V1 and V1 to V2. For these two types of predictions we followed a very similar method. For the stimulus to V1 mappings, we had an existing model with which we could compare. We used this to reinforce the confidence in the results of our method, as there is no existing model for the V1 to V2 mappings.

Estimation of error:

In order to measure the correctness of our predictions, we chose to use the average of the sum of squared errors. This method allows us to compare the predictions of each method in order to know how they perform and to what extent their results are reliable. The numbers given in Table 1 have been obtained with the predictions on the test set (two hundreds time points per voxel) and have also been averaged on several different voxels. These numbers have to be analyzed together with prediction curves and the mapping.

Stimulus to V1:

Results for the prediction of an LV1 voxel's activation from the visual stimulus can be found in the two first columns of Table 1. It shows that these predictions are quite accurate. The SVM predictions tend to outperform LASSO predictions even if they are close for the full stimulus. Fig. 2(a) is a good error analysis example of why the error of LASSO with PCA is significantly higher, while the mapping is still correct: it appears that it tends to amplify the actual variations of the signal. However the way the signal varies is still accurate. SVM predictions lose less accuracy when non crucial information is removed.

These results give us the possibility to detect the receptive field of each voxel (Fig. 1). This part is particularly important as we can then compare our results with the receptive fields given by Professor Wandell's Lab. The strong correspondence between our results and the lab's results reinforces the accuracy of the prediction and indicates that the method we followed is efficiently solving the problem.

V1 to V2:

Results for V1 to V2 predictions can be found in the third column of the table. These results are for our prediction with the 600 principal components of V1, but we tried several numbers and the results seem very robust to this parameter.

Again, SVM predictions are slightly better than LASSO predictions, but they are very close. These errors are low and show a very accurate prediction of V2 activation with V1 information. This can also be seen in Fig 2

prediction curves on which one can see that predictions are very rarely far from the actual signal.

It also appears that we are able to predict far more precisely V2 activation from V1 than V1 activation from the signal. This result was absolutely not obvious a-priori as the data we have on the stimulus signal are perfect, but fmri data of V1 activation are far from this resolution. One possible explanation for the lower results for the stimulus to V1 predictions could be due to what is determined to be the true stimulus. We assume that the stimulus is limited only to the image that the subject sees, but there may be other factors that contribute to what the patient actually sees.

These good results, as well as the confidence we can have in the method as a result of verifying the stimulus to V1 predictions with a given model, enable us to give an accurate activation mapping for V2 voxels. This mapping which can be seen in Fig. 3c achieves our main goal.

Conclusion

We were able to build successful models for both stimulus to V1 prediction and V1 to V2 prediction. By implementing machine learning algorithms, we were able to develop accurate models. These results will be useful to neuroscientists in understanding the relations between these two regions of the visual cortex. These methods could be generalized to determine mappings in other regions of the Visual Cortex as well as elsewhere in the brain.

Acknowledgments

We acknowledge Jonathan Winawer and Brian Wandell from the Psychology department at Stanford University for providing the data, comparative model, computing power and guiding us through this project.

References

[1] Brian A. Wandell, “Foundations of Vision”, pp. 187-191.

[2] Tibshirani, R. (1996). “Regression shrinkage and selection via the lasso”, J. R. Statis. Soc. B., Vol. 58, No. 1, pp. 267-288.

[3] http://cran.r-project.org/