Feature Induction Using Boosting and Logistic Regression on fMRI Images

Melissa K. Carroll[1], Miroslav Dudík[1], Robert E. Schapire[1], Kenneth A. Norman[2]
Departments of Computer Science[1] and Psychology[2]
Princeton University
Princeton NJ 08540
{mkc,mdudik,schapire,knorman}@princeton.edu

Abstract

Early efforts in fMRI classification were limited in that individual voxels were used as features (e.g. [1]), yet voxels divide images into regions that do not directly correspond to underlying neural activity. A growing trend is to perform spatial smoothing that captures the correlation between nearby voxels. Unfortunately, the optimal spatial resolution for this smoothing is unknown and likely varies across brain regions and cognitive tasks. The present work describes two methods that induce features of varying size and shape and use them to produce additive models that offer the potential for easy interpretability.

1 Methods

fMRI images were collected while two subjects viewed objects from among 7 classes, similar to those used in [2], with the goal of classifying the fMRI images by the class of object being viewed. 2D cortical surface mapped data were used.

The first method involves computing features similar to those in the Viola and Jones [3] algorithm for 2D object detection (Figure 1) and producing a model from them using AdaBoost with stumps [4].

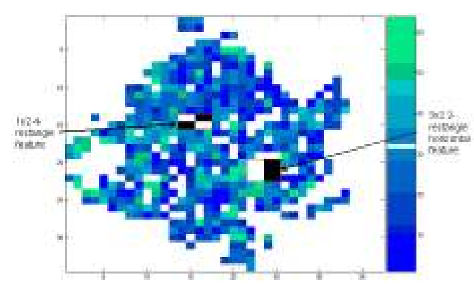

Figure 1. Sample Viola-Jones features on a brain image. Total values in white regions are subtracted from total values in gray regions. Types include 2-,3-, and 4-rectangle features, as well as 1-rectangle features, which merely sum over regions. Features for all permutations of type, size, and position are generated.

The second method uses a feature induction technique in logistic regression to search for optimal features that are free to vary by both size and shape. To model the data, a set of weights is chosen that maximizes the log likelihood and l_1-style penalty is used to prevent overfitting. A sequential boosting-like update algorithm along the lines of [5] and [6] is used, which in each iteration picks a single feature and class label and greedily updates the associated weight. In one version, only features derived from single pixels are considered. The second version begins with single-pixel features, but in each iteration, after updating the weight, new features are added, which are obtained by adding the value of a single adjacent pixel to the last updated region. This approach is similar to the one used by [7] in the context of natural language modeling.

2 Results

Both methods were evaluated using leave-one-run-out cross-validation. Figure 2 shows the classification performance of the first method. Given the large feature sets that result from calculating all of the possible features for an image, the total feature set is randomly sampled; the percent of features included in the model is varied.

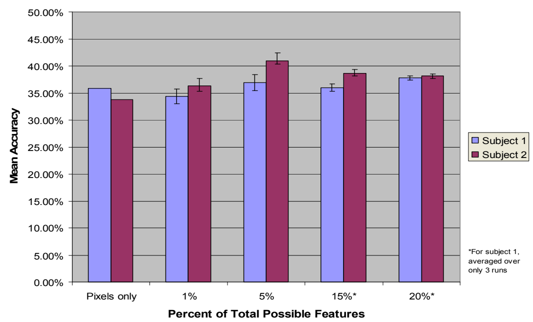

Figure 2: Classifier performance (averaged over 5 runs when using randomization).

Figure 3 demonstrates the potential of the second method for producing interpretable results. Not only are such graphs easily produced, but the grown-features variant can produce less noisy images than does the single-pixel-feature variant.

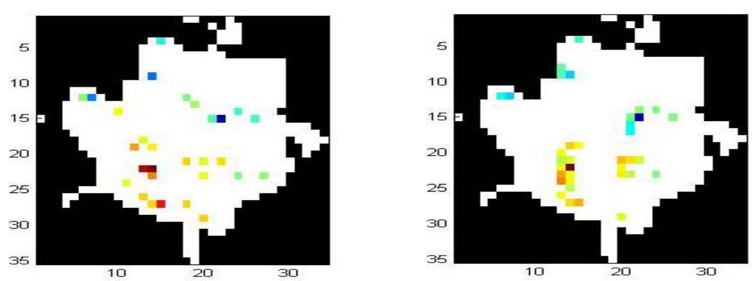

Figure 3: Maps of one subject's VT region for the "dog face" model in which blue pixels reflect lower weights and red pixels higher weights. The single-pixel-feature model is shown on the left and the grown-feature model is on the right.

3 Conclusion

In some cases, classification accuracy can be improved by producing linear models of induced features that capture spatial correlation. While not sacrificing performance, these methods also offer the promise of easily interpreted models.

References

[1] Mitchell, T. M., Hutchinson, R., Niculescu, R. S., Pereira, F., Wang, X., Just, M., & Newman, S. (2004). Learning to Decode Cognitive States from Brain Images. *Machine Learning* **57** (1-2), pp. 145–175.

[2] Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai,, A., Schouten, J.L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* **293**, pp. 2425-2429.

[3] Viola, P. & Jones, M. (2001). Rapid object detection using a boosted cascade of simple features. *CVPR*.

[4] Freund, Y. & Schapire, R. E. (1996). Experiments with a New Boosting Algorithm. *Proceedings of the Thirteenth International Conference on Machine Learning,* pp. 148-156. Morgan Kaufmann.

[5] Collins, M., Schapire, R.E., & Singer, Y. (2000). Logistic regression, AdaBoost and Bregman distances. *Proceedings of the Thirteenth Annual Conference on Computational Learning Theory*.

[6] Dudík, M., Phillips, S.J., & Schapire, R.E. (2004). Performance guarantees for regularized maximum entropy density estimation. *Proceedings of the Seventeenth Annual Conference on Learning Theory* pp. 472-486. Springer-Verlag.

[7] Della Pietra, S., Della Pietra, V. & Lafferty, J. (1997) Inducing features of random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **19**(4), pp. 1-13.