Published in 2012 15th International Conference on Information Fusion, FUSION: Singapore, 9-12 July: Proceedings. pp. 593-599.

Clustered subsampling for clinically informed diagnostic brain mapping

Malin Björnsdotter*, Diego Sona†, Sonny Rosenthal‡ and Justin Dauwels*
*School of Electrical and Electronics Engineering, Nanyang Technological University, Singapore; MBjornsdotter@ntu.edu.sg
†Pattern Analysis and Computer Vision, Fondazione Istituto Italiano di Tecnologia, Italy
‡Wee Kim Wee School of Communication and Information, Nanyang Technological University, Singapore

***Abstract*—Brain based diagnostic systems have recently received attention as a tool in the characterization and diagnosis of a variety of neurodevelopmental and psychiatric disorders. Nonetheless, a majority of disorders are still diagnosed entirely based on symptom assessments and behavioral correlates. We therefore propose a method for fusing brain responses with clinical measures for improved diagnosis. To this end, we utilized the flexibility of clustered random subspace brain mapping to detect regions where brain responses in conjunction with a clinical measure could reliably differentiate patients from control subjects. We demonstrate the approach on realistically simulated functional magnetic resonance imaging (fMRI) brain activity and a clinical parameter. We show that the method efficiently identifies brain regions where fused analysis of brain responses and clinical parameters improves diagnosis compared to either measure alone. The proposed method is easy to implement and highly flexible, offering an appealing basis for multimodal brain mapping.**

I. Introduction

The contemporary neuroscientist has access to a plethora of increasingly sophisticated measuring techniques reflecting the wide range of systems, levels and modalities found in biological organisms – from genes through single spikes to distributed brain activity and various measures of behavior. A complete understanding of the human nervous system in health and disease can clearly not be achieved from study of a single level or modality, and the research community is therefore turning their attention to the possibility of designing studies that span multiple modalities (see e.g. [1]).

Concurrent with the development of increasingly advanced neuroimaging technologies, there has been a surge of studies investigating the utility of brain based diagnosis (for disorders such as autism [2], depression [3], attention deficit disorder [4] and Alzheimer's disease [5]). Nonetheless, a majority of disorders are still preferentially diagnosed based on clinical symptom assessments or behavioral correlates, and few rely mainly on brain based diagnostics (see e.g. [6], [7]). Instead, it appears that an efficient neuroimaging diagnostic aid should integrate multiple modalities, including clinical symptom assessments and, where relevant, behavioral correlates.

A consequence of the long history of relative fragmentation of neuroscientific research, however, is that each modality is associated with its own peculiar processing and analysis approaches. Conventional cross-modal studies process data independently and, subsequently, estimate inter-modality coupling using post hoc juxtaposition through statistical correlation measures. Truly integrated analysis, in contrast, makes explicit a priori assumptions about the coupling between the modalities. The superiority of such fused analysis over juxtaposition techniques was recently demonstrated in a study of the coupling between haemodynamic (measured through functional magnetic resonance imaging; fMRI) and electrophysiological (acquired using electroencephalography; EEG) brain activity: from an information-theory point of view, the distribution of responses from the fused modalities was more informative than any single modality [8]. Consequently, truly synergistic multimodal analysis is a rapidly growing focus of interdisciplinary research. Moreover, conventional correlation approaches are poorly suited for generalizing the obtained model to new data. Thus, integrated analyses should be both highly sought after and highly beneficial in clinical situations, e.g., for disease diagnosis.

Addressing these issues, we propose a method for direct cross-modal neuroimaging analysis, where brain responses and clinical scores are fused. To this end, we utilized machine learning – a domain that provides highly appealing tools for sophisticated recognition of otherwise unintelligible pattern representations embedded in complex data structures [9]. Machine learning is generic, flexible (e.g. classification, regression, clustering), powerful (e.g. multivariate, linear and nonlinear), predictive, and applicable to various data modalities with minor modality-specific preprocessing. Importantly, machine learning is data-driven, taking advantage of actual data examples to capture relevant characteristics without requiring prior information on underlying probability distributions. The appealing properties of machine learning has attracted considerable attention in various sectors of the community in recent years, including analysis of fMRI [10].

We focused on fMRI-quantified brain activity, which is a safe, non-invasive technique for exploring function-structure relationships based on changes in local blood delivery within different regions of the brain. The fMRI signal is derived from magnetic properties of deoxygenated blood, reflecting temporal changes in activity as the brain processes information. Such processing is inherently distributed and multivariate, and, therefore, recent applications of machine learning has had crucial impact on fMRI analysis methods [10]. However, fMRI yields a staggering amount of data - hundreds of thousands of measurements, called voxels - which calls for efficient feature selection. We therefore build on a locally multivariate brain

mapping method based on random subspace sampling that we recently proposed [11], [12]. We have previously demonstrated the utility of this method for efficient and sensitive multivariate brain mapping both on simulated data [12] and authentic neuroscientific problems [13], [14].

Random subspace sampling maps are constructed by randomly selecting a number of voxel subsets, each on which a classifier is trained. The random subspace ensemble approach is intuitive and simple, and yields highly competitive results compared to other voxel selection methods [15], [16]. Notably, in addition to providing an appealing tool for mapping brain regions, random subsampling is highly flexible and allows easy incorporation of external parameters as well as usage of any information measure (including non-linear classifiers). In this study, we therefore utilized the flexibility, efficiency and mapping sensitivity of clustered subspace brain mapping for multimodal neuroimaging analysis. Specifically, we investigated fused analysis of brain activity and an external parameter, here interpreted as a clinical measure, to propose a method that identifies brain areas where synergistic multimodal analysis improves diagnostic accuracy.

We demonstrate the utility of the method on realistically simulated data. We modeled the fMRI data to emulate brain scans of 64 individuals, half of whom are patients with a condition or disorder that yields deviating activity in a specific brain region. This region corresponds to brain areas associated with impaired processing in individuals diagnosed with the disorder, such as impairment of the right posterior temporal sulcus in autism [17]. The simulated clinical parameter emulates a quantitative assessment assigned to each individual, and is inspired by the social responsiveness score (SRS) in autism [18]. Although the SRS is indicative of autism, it is far from perfect: studies report that the SRS has a sensitivity of 0.74 to 0.80 and a specificity of 0.69 to 1.00 in detecting individuals diagnosed with autism [19].

Specifically, the simulated data aimed at emulating a situation where patients are reliably distinguished from controls only when clinical assessment parameters and region-specific brain response are analyzed simultaneously. Consequently, the objective of the proposed algorithm was to efficiently identify brain regions in which fused analysis is more diagnostic than either of the measures alone.

II. Methods

Simulated data

Clinical assessment score: The clinical assessment score was simulated as one value per individual. All subjects' values were randomly generated uniformly in the range 1:10, after which patients' values were adjusted up by 4 points. Hence, patients' values were poorly separated from controls based on this simulated behavioral correlate alone (Figure 1).

Brain activity: The brain activity was based on simulated data from a previous study [20]. In summary, the data was simulated in a cortex mask, obtained from a real subject, containing 28502 voxels. Data modeling followed a block-design with two conditions for a functional voxel resolution

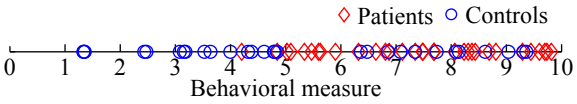

Fig. 1. The simulated clinical assessment score.

to 3×3×3 mm with 64 stimulations (equal number of each condition) lasting 10.5 s each. Blood-oxygen-level dependent responses were simulated by convolving the stimuli with a double-gamma hemodynamic response function. Voxelwise activations were added to temporally autocorrelated noise, modeled as follows:

$$R_a(t) = \rho_k R_0(t-1) + \sqrt{1-\rho_k^2} R_0(t) \quad (1)$$

where R_0 is random Gaussian noise and $\rho_k : N(0.5, 0.1)$ allowed adjustment of the amount of autocorrelation at voxel k. Single-trial responses were estimated by fitting a voxel-wise general linear model (GLM) with one predictor representing each single stimulus (obtained by convolution of a boxcar with a double-gamma hemodynamic response function; [21]) and one linear predictor accounting for a within-trial linear trend. The regressor β coefficients were subsequently obtained as an estimate of each trial response.

Here, the simulated data was re-interpreted such that each sample represents an individual's brain response, and each condition either a control subject or a patient. The final simulated data hence consisted of 64 individuals, 32 of whom were designated as patients and half as control subjects.

In one brain region (142 voxels), activity differentiating control subjects from patients was added (indicated in blue in Figure 2). These voxels were randomly assigned to one of two populations (patients > controls; controls > patients). The contrast-to-noise ratio (CNR; the difference between the group responses compared to the standard deviation of the noise) was set 0.2, ensuring that a multivariate analysis (i.e. joint analysis of multiple voxels) was required to reliably detect differences. In addition, all subject's brain responses were adjusted such that the patients could be clearly separated from controls only in fusion with the simulated clinical parameter (Figure 3).

Clustered subsampling brain mapping

Kuncheva et al. proposed the following method of constructing a random subspace ensemble for brain decoding: let $\mathcal{V} = v_1, ..., v_n$ be the set of n voxels. L voxel subsets, each of size N, are drawn without replacement from a uniform distribution over V and a classifier model $\mathcal{M}$ is trained on each subset [16]. The ensemble decision is subsequently made by majority vote among the L classifiers. As shown in Algorithm 1, random subspace maps are constructed in a similar fashion [11], [12]. A function Φ computes the classification accuracy of model $\mathcal{M}$, which is trained on each voxel subset, and the accuracy is added to each voxel (and stored in vector R). Importantly, however, the voxel subsets are selected such that they form an approximately spherical cluster: a voxel v is

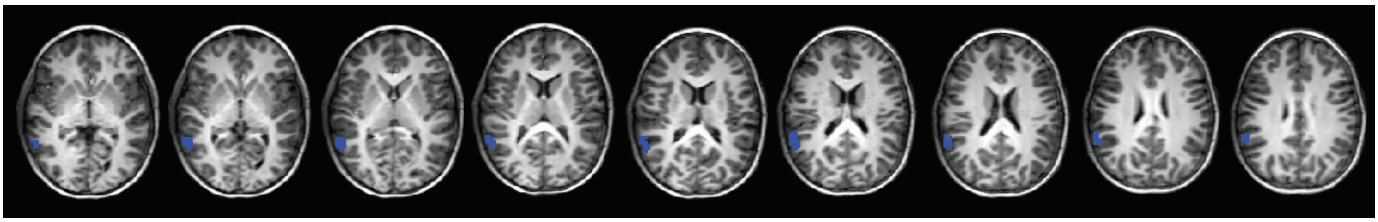

Fig. 2. The brain region containing activity differentiating control subjects from patients, indicated in blue, across a number of axial brain slices.

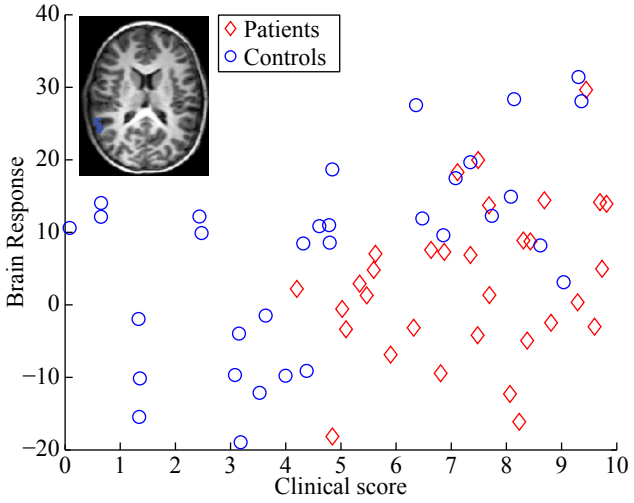

Fig. 3. Simulated brain responses from a randomly selected voxel in the diagnostic region (shown in blue) as a function of the clinical assessment score.

randomly selected in the brain volume and voxels neighboring v are included in the subset. The process is iterated I times, after which final voxel-wise values are obtained by averaging across the values obtained in each iteration (i.e. normalization of R; Figure 4). The resulting value reflects the usefulness of each voxel in multivariate combination with other voxels: the more informative the voxel, the more likely subsets containing the voxel will obtain high results which will boost the average value for that voxel. The required number of times the model must be trained is $L = \lfloor \frac{n}{N} \rfloor$ per iteration, and, in the current implementation, there is no majority vote to obtain a single decoding accuracy representative of the brain volume.

Algorithm 1 Clustered sampling brain mapping

Input: Dataset X, voxel set $\mathcal{V}$ and learning model $\mathcal{M}$
- Initialize vector R=[0, . . . , 0]
repeat
  while: $\mathcal{V}$ is not empty **do**
    - Randomly select a voxel $v \in \mathcal{V}$
    - Select the voxels $\mathcal{V}^k \subset \mathcal{V}$ in the neighborhood of v
    - Train the model $\mathcal{M}$ on voxels $\mathcal{V}^k$
    - $\forall v_i \in \mathcal{V}^k$ update the rating $R_{v_i} = R_{v_i}$ + $\Phi(\mathcal{M})$
    - Remove $\mathcal{V}^k$ from $\mathcal{V}$
  end while
until convergence
- Normalize R
return R

In the current study, Algorithm 1 is modified to include the clinical parameter (variable $\mathcal{C}$) in the modelling phase (Algorithm 2). Here, each time local brain responses are input to the model $\mathcal{M}$, they are paired with the clinical variable such that both brain responses and variable are used to train $\mathcal{M}$ and in function Φ_{v_i} to evaluate the classification accuracy.

Algorithm 2 Modified clustered sampling brain mapping

Input: Dataset X, voxel set $\mathcal{V}$, variable $\mathcal{C}$ and learning model $\mathcal{M}$
- Initialize vector R=[0, . . . , 0]
repeat
  while: $\mathcal{V}$ is not empty **do**
    - Randomly select a voxel $v \in \mathcal{V}$
    - Select the voxels $\mathcal{V}^k \subset \mathcal{V}$ in the neighborhood of v
    - Train the model $\mathcal{M}$ on voxels $\mathcal{V}^k$ and variable $\mathcal{C}$
    - $\forall v_i \in \mathcal{V}^k$ update the rating $R_{v_i} = R_{v_i}$ + $\Phi(\mathcal{M})$
    - Remove $\mathcal{V}^k$ from $\mathcal{V}$
  end while
until convergence
- Normalize R
return R

Parameters

Two parameters require specification: the number of voxels in the subsets (i.e., search sphere volume) and the number of times the algorithm is iterated (i.e., the number of times the model is trained on each voxel). In the current study, we fixed the search sphere volume size to a radius of 3 voxels and iterated the algorithm 5 times. The algorithm was implemented in Matlab (The Mathworks, Massachusetts, USA).

Measure of diagnostic performance

As a learning model (i.e. to quantify the diagnostic power of the voxel subsets) we used five-fold cross-validation scheme with a support vector machine (SVM, in the LS-SVMlab implementation; http://www.esat.kuleuven.be/sista/lssvmlab/ [22]) with a linear kernel.

We computed voxel-wise significance levels in the average map using the binomial distribution. The obtained p-values were corrected for multiple comparisons through controlling for the false discovery rate (FDR) such that q<0.05 [23]. Also, the voxel subset with the highest classification accuracy was identified and this peak classification accuracy was reported.

As a measure of the mapping performance (i.e. the ability of the mapping scheme to identify the simulated diagnostic voxels), the receiver operating characteristic (ROC) curve (a plot of the voxel detection sensitivity versus 1-specificity for varying map thresholds) was computed for each map and the area under the curve (AUC), was obtained. A value of 1

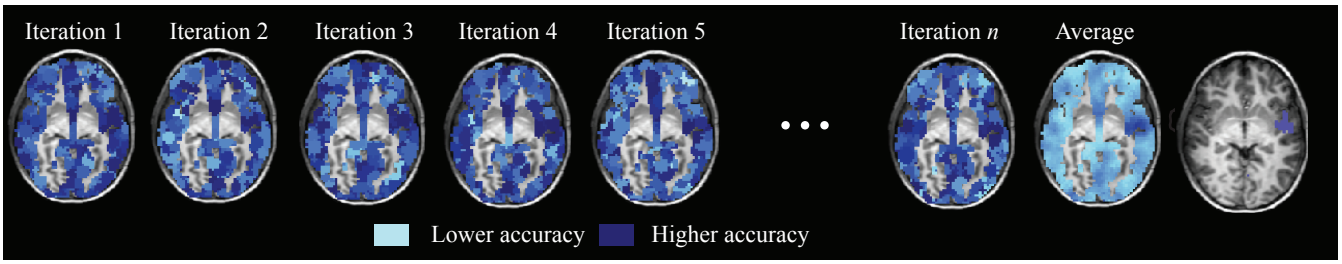

Fig. 4. Clustered subsampling brain mapping maps are constructed by iteratively sampling spherical volumes within which a classification accuracy is computed. The mean across these volumes is then computed to produce a map showing the average contribution of each voxel. The region with true diagnostic (high-accuracy) voxels in this example is indicated in blue in the right-most panel.

corresponds to a perfect map, whereas the chance value in the current dataset was 0.9 (due to the large number of true negatives). Also, we noted the computing time required to obtain the results as well as the number of times the classifier model was estimated (i.e. trained) and evaluated.

Comparison with alternate methods

For comparison, we repeated the five-fold cross-validation scheme in four additional ways: on the clinical variable only, on single voxel brain responses only (i.e. on the data from each individual voxel independently), on multivariate brain responses (i.e. Algorithm 1), and, finally, on single voxel brain responses and the clinical variable combined. Maps from the single voxel analysis were obtained by saving the voxel-wise classification accuracies. Again, we computed the area under the ROC curve and noted the peak classification performance, computing time and number of times the classifier model was estimated and evaluated.

III. Results

The results are summarized in Table I.

Detection of brain voxels: The clustered subsampling approach successfully identified the brain region where fused brain and behavior analysis could separate patients from controls (Figure 5A): the area under the ROC was 0.999 . The peak search volume was also located in this region (Figure 5B). The single-voxel brain response mapping was, on the contrary, unsuccessful in detecting the diagnostic voxels, with an area under the ROC curve of 0.888 and suffering from a severe number of false positives (Figure 6A). The performance improved only marginally after fusion with the clinical parameter to 0.900 (Figure 6C). The multivoxel brain mapping approach, on the other hand, was second most effective in detecting the diagnostic voxels with an area under the ROC curve of 0.993 (Figure 6B).

Diagnostic accuracy: The maximum classification accuracy in the peak search volume (show in in Figure 5B) was 90.3%. Contrasting this result, the simulated clinical measure alone yielded a classification accuracy of 67.0%, whereas the multivoxel brain response obtained an accuracy of 81.4%. The single-voxel brain response alone and the single-voxel response fused with the clinical variable achieved relatively high classification accuracies, at 79.9% and 86.2%, respectively. However, it should be noted that the latter two are likely to reflect an exaggerated diagnostic power given the large number of computations involved (28502, compared to only 3810 for the multivoxel brain mapping approaches).

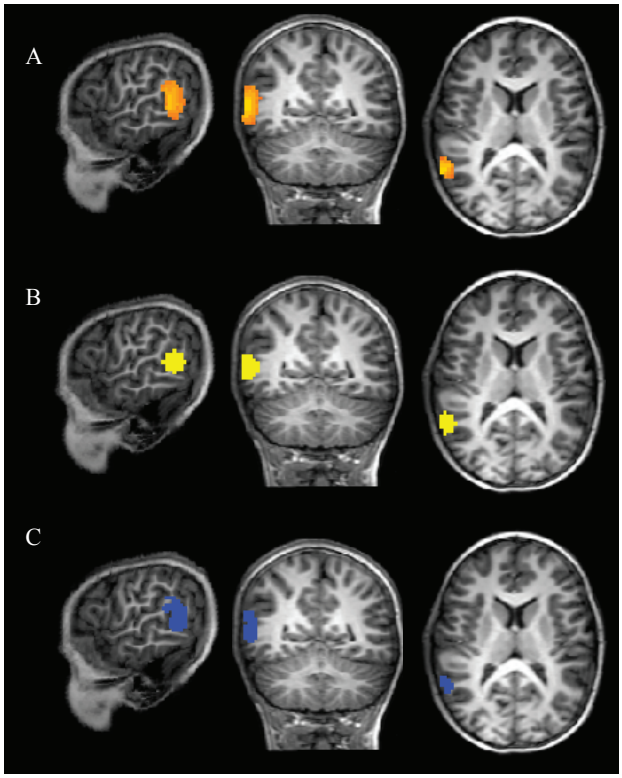

Fig. 5. A) Voxels in the average classification map surviving the $p < 0.05$ corrected threshold. B) The voxel subset with the highest classification. C) The true simulated diagnostic voxels.

Computational efficiency: The multivoxel algorithms, including the proposed method, used 3810 computations (i.e. training and testing of the classifier) to compute the results, which required 18.6 minutes. The single-voxel methods were faster at 12.8 and 14.1 minutes, but required the classifier model to be estimated and evaluated 28502 times (i.e. as many computations as there were voxels in the brain).

TABLE I
RESULTS

Method	Peak classification (%)	Mapping performance (AUC)	Time (minutes)	Nr. computations
Clinical score	67.0	-	-	1
Single voxel brain response	79.9	0.888	12.8	28502
Multivoxel brain response	81.4	0.993	18.6	3810
Single voxel fused brain and clinical score	86.2	0.900	14.1	28502
Multivoxel fused brain and clinical score	90.3	0.999	18.6	3810

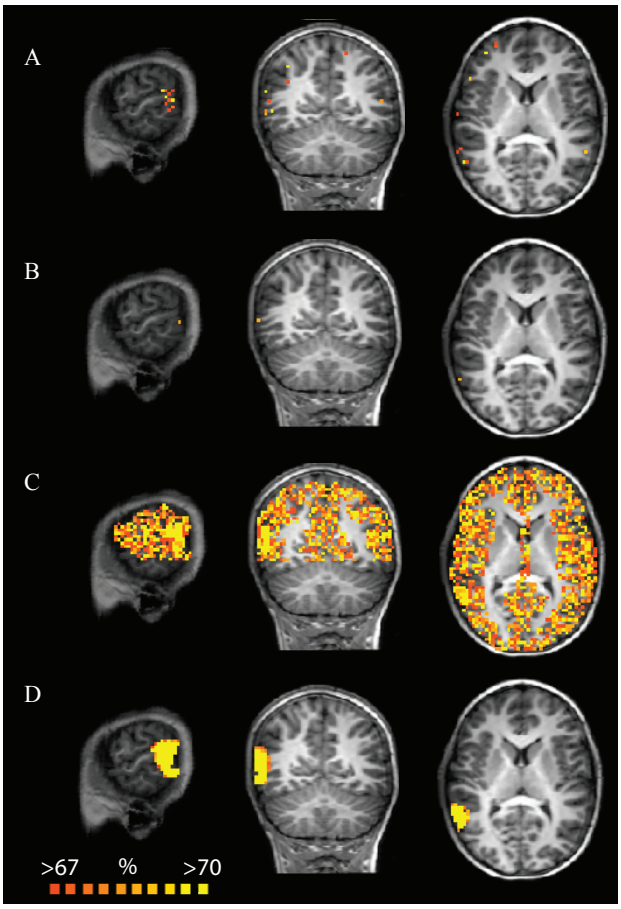

Fig. 6. Results from all compared brain mapping methods: A) Single voxel brain response B) Multivoxel brain response C) Single voxel fused brain and behavior D) Multivoxel fused brain and behavior. The maps are thresholded at the classification accuracy for behavior alone (67%).

IV. DISCUSSION

We have presented a computationally efficient multivariate brain mapping approach which exploits fused classification of brain responses and clinical scores for improved neurodiagnosis. We demonstrated that fused analysis is more effective in brain mapping and, hence, achieves higher diagnostic accuracy, than either measure alone, on simulated data where either measure were only partially diagnostic.

The clinical score was modeled similar to the social responsiveness score used as a diagnostic aid in autism [18], and the simulated brain activity was designed to mimic realistic fMRI responses. Moreover, the brain data was simulated to be inherently multivariate such that conventional univariate statistics, including the general linear model (GLM [21]), would fail to identify the diagnostic voxels. Multiple studies have demonstrated the occurrence of subtle brain activity changes encoded across multiple voxels which conventional univariate statistics fail to capture [10], [24]–[30]. We therefore believe that the proposed method should work well with authentic data, in particular in situations where patient's brain responses represent a complex impairment.

We compared the proposed algorithm with single-voxel methods, which, albeit being faster (at 12.8 and 14.1 compared to 18.6 minutes), were substantially worse at detecting the diagnostic voxels (with areas under the ROC curve of around 0.9, i.e. chance level, compared with a nearly perfect map for the proposed method). Although the obtained maps are clearly superior in the fused multivoxel approach, which should justify the increased time requirements in some contexts, the achieved classification accuracy for the single-voxel fused brain and behavior method (86.2%) rivaled the proposed algorithm (90.3%). Nevertheless, the single-voxel approach estimated the diagnostic accuracy 28502 times, compared to a mere 3810 times for the multivoxel methods, and this high classification accuracy likely reflects overfitting.

The proposed algorithm shares many similarities with the Searchlight algorithm [31], which is highly popular for multivariate fMRI brain mapping (see e.g. [32]–[37]). Nonetheless, our clustered random subsampling algorithm requires dramatically less computation times (in the order of 75% less) for the same mapping sensitivity [12] and is therefore a more appealing alternative in, for example, large patient studies.

In the current implementation, we only considered locally multivariate brain responses. However, improved prediction power may be achieved from combing multiple brain regions. This may be particularly useful when considering spatially disjoint brain impairments, such as in autism where connectivity is disrupted (by e.g. considering the posterior superior temporal sulcus and the fusiform gyrus, both part of an extended network that has been implied as being disrupted in autism [17]). More sophisticated approaches where fused brain-behavior analysis is feasible is through Memetic algorithms [38] or particle swarm optimization [39]. Also, we simply reported the classification accuracy from the voxel subset with the best performance. An ensemble decision obtained majority vote among the classifiers, akin to what was implemented by Kuncheva and colleagues [16], is likely to improve the accuracy.

A major benefit of the proposed method is that it can be applied in conjunction with any modality and in combination

with a large variety of information measures with minor adjustment: it can be cross-modally used in conjunction with classification, including both linear and non-linear classifiers on categorical data, but also with regression (using e.g. relevance vector machines, RVMs) to predict continuous variables (such as the severity of a disorder). Other modalities may include behavioral correlates, electroencephalography (EEG), physiological parameters and so on. In particular, the method has potential in the domain of imaging genomics: a wide range of cognitive disorders have been studied in the framework of imaging genomics, including psychiatric illness (such as e.g. schizophrenia; [40]) and affective disorders (including depression [41], [42] and bipolar disorder [43]). Genomics and brain imaging belong to historically disparate fields of research, and contemporary genetic linkage studies are therefore typically limited to use mass-univariate linear modeling and post hoc juxtaposition correlation approaches to examine gene-phenotype relationships. The proposed algorithm has decided potential in bridging this gap.

V. Conclusion

We have presented a method for mapping brain regions where neural responses in combination with an external variable, here modeled as a clinical correlate, can predict the diagnostic category of individuals. The simplicity and flexibility of the method makes it an appealing starting-point for a general-purpose cross-modal brain mapping tool.

Acknowledgment

M.B. was supported by the Wenner-Gren Foundations and the Marie Curie International Outgoing Fellowship.

References

[1] R. J. Perrin, A. M. Fagan, and D. M. Holtzman, "Multimodal techniques for diagnosis and prognosis of alzheimer's disease." *Nature*, vol. 461, no. 7266, pp. 916–922, Oct 2009. [Online]. Available: http://dx.doi.org/10.1038/nature08538

[2] M. D. Kaiser and K. A. Pelphrey, "Disrupted action perception in autism: Behavioral evidence, neuroendophenotypes, and diagnostic utility," *Developmental Cognitive Neuroscience*, vol. 2, no. 1, p. 2535, 2010.

[3] X. Zhang, Z. S. Yaseen, I. I. Galynker, J. Hirsch, and A. Winston, "Can depression be diagnosed by response to mother's face? a personalized attachment-based paradigm for diagnostic fmri." *PLoS One*, vol. 6, no. 12, p. e27253, 2011. [Online]. Available: http://dx.doi.org/10.1371/journal.pone.0027253

[4] C. Z. Zhu, Y. F. Zang, M. Liang, L. X. Tian, Y. He, X. B. Li, M. Q. Sui, Y. F. Wang, and T. Z. Jiang, "Discriminative analysis of brain function at resting-state for attention-deficit/hyperactivity disorder." *Med Image Comput Comput Assist Interv*, vol. 8, no. Pt 2, pp. 468–475, 2005.

[5] E. E. Tripoliti, D. I. Fotiadis, M. Argyropoulou, and G. Manis, "A six stage approach for the diagnosis of the alzheimer's disease based on fMRI data." *J Biomed Inform*, vol. 43, no. 2, pp. 307–320, Apr 2010. [Online]. Available: http://dx.doi.org/10.1016/j.jbi.2009.10.004

[6] B. P. S. I. Walsh P, Elsabbagh M, "In search of biomarkers for autism: scientific, social and ethical challenges." *Nat Rev Neurosci.*, vol. 12, no. 10, pp. 603–12., 2011.

[7] J. L. Stevenson and K. A. Kellett, "Can magnetic resonance imaging aid diagnosis of the autism spectrum?" *J Neurosci*, vol. 30, no. 50, pp. 16 763–16 765, Dec 2010. [Online]. Available: http://dx.doi.org/10.1523/JNEUROSCI.4946-10.2010

[8] D. Ostwald, C. Porcaro, and A. P. Bagshaw, "An information theoretic approach to eeg-fmri integration of visually evoked responses." *Neuroimage*, vol. 49, no. 1, pp. 498–516, Jan 2010. [Online]. Available: http://dx.doi.org/10.1016/j.neuroimage.2009.07.038

[9] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*. Wiley-Interscience Publication, 2000.

[10] J. Haynes and G. Rees, "Decoding mental states from brain activity in humans," *Nature Reviews Neuroscience*, vol. 7, no. 7, pp. 523–534, July 2006. [Online]. Available: http://dx.doi.org/10.1038/nrn1931

[11] D. Sona and P. Avesani, "Feature rating by random subspaces for functional brain mapping," in *Brain Informatics*, 2010, pp. 112–123.

[12] M. Björnsdotter, K. Rylander, and J. Wessberg, "A monte carlo method for locally multivariate brain mapping." *Neuroimage*, vol. 56, no. 2, pp. 508–516, May 2011. [Online]. Available: http://dx.doi.org/10.1016/j.neuroimage.2010.07.044

[13] D. Sona and P. Avesani, "Multivariate brain mapping by random subspaces," vol. 0. Los Alamitos, CA, USA: IEEE Computer Society, 2010, pp. 2576–2579.

[14] V. I. Petkova, M. Björnsdotter, G. Gentile, T. Jonsson, T.-Q. Li, and H. H. Ehrsson, "From part- to whole-body ownership in the multisensory brain." *Curr Biol*, vol. 21, no. 13, pp. 1118–1122, Jul 2011. [Online]. Available: http://dx.doi.org/10.1016/j.cub.2011.05.022

[15] L. I. Kuncheva and J. J. Rodrguez, "Classifier ensembles for fMRI data analysis: an experiment." *Magn Reson Imaging*, vol. 28, no. 4, pp. 583–593, May 2010. [Online]. Available: http://dx.doi.org/10.1016/j.mri.2009.12.021

[16] L. I. Kuncheva, J. J. Rodriguez, C. O. Plumpton, D. E. J. Linden, and S. J. Johnston, "Random subspace ensembles for fMRI classification." *IEEE Trans Med Imaging*, vol. 29, no. 2, pp. 531–542, Feb 2010. [Online]. Available: http://dx.doi.org/10.1109/TMI.2009.2037756

[17] M. D. Kaiser, C. M. Hudac, S. Shultz, S. M. Lee, C. Cheung, A. M. Berken, B. Deen, N. B. Pitskel, D. R. Sugrue, A. C. Voos, C. A. Saulnier, P. Ventola, J. M. Wolf, A. Klin, B. C. Vander Wyk, and K. A. Pelphrey, "Neural signatures of autism." *Proc Natl Acad Sci U S A*, vol. 107, no. 49, pp. 21 223–21 228, Dec 2010. [Online]. Available: http://dx.doi.org/10.1073/pnas.1010412107

[18] J. N. Constantino and R. D. Todd, "Autistic traits in the general population: a twin study." *Arch Gen Psychiatry*, vol. 60, no. 5, pp. 524–530, May 2003. [Online]. Available: http://dx.doi.org/10.1001/archpsyc.60.5.524

[19] S. Bölte, E. Westerwald, M. Holtmann, C. Freitag, and F. Poustka, "Autistic traits and autism spectrum disorders: the clinical validity of two measures presuming a continuum of social communication skills." *J Autism Dev Disord*, vol. 41, no. 1, pp. 66–72, Jan 2011. [Online]. Available: http://dx.doi.org/10.1007/s10803-010-1024-9

[20] M. Björnsdotter, K. Rylander, and J. Wessberg, "A monte carlo method for locally multivariate brain mapping." *Neuroimage*, Jul 2010. [Online]. Available: http://dx.doi.org/10.1016/j.neuroimage.2010.07.044

[21] K. J. Friston, A. P. Holmes, K. J. Worsley, J. P. Poline, C. D. Frith, and R. S. J. Frackowiak, "Statistical parametric maps in functional imaging: A general linear approach," *Human Brain Mapping*, vol. 2, no. 4, pp. 189–210, 1994.

[22] J. Suykens, T. V. Gestel, J. D. Brabanter, B. D. Moor, and J. Vandewalle, *Least Squares Support Vector Machines*. World Scientific, 2002.

[23] C. R. Genovese, N. A. Lazar, and T. Nichols, "Thresholding of statistical maps in functional neuroimaging using the false discovery rate." *Neuroimage*, vol. 15, no. 4, pp. 870–878, Apr 2002. [Online]. Available: http://dx.doi.org/10.1006/nimg.2001.1037

[24] J. V. Haxby, M. I. Gobbini, M. L. Furey, A. Ishai, J. L. Schouten, and P. Pietrini, "Distributed and overlapping representations of faces and objects in ventral temporal cortex." *Science*, vol. 293, pp. 2425–2430, 2001.

[25] D. D. Cox and R. L. Savoy, "Functional magnetic resonance imaging (fMRI) 'brain reading': detecting and classifying distributed patterns of fMRI activity in human visual cortex." *Neuroimage*, vol. 19, no. 2 Pt 1, pp. 261–270, June 2003.

[26] A. J. O'toole, F. Jiang, H. Abdi, and J. V. Haxby, "Partially distributed representations of objects and faces in ventral temporal cortex," *Journal of Cognitive Neuroscience*, vol. 17, no. 4, pp. 580–590, 2005.

[27] S. M. Polyn, V. S. Natu, J. D. Cohen, and K. A. Norman, "Category-specific cortical activity precedes retrieval during memory search," *Science*, vol. 310, no. 5756, pp. 1963–6, 2005.

[28] Y. Kamitani and F. Tong, "Decoding the visual and subjective contents of the human brain," *Nature Neuroscience*, vol. 8, no. 5, pp. 679–685, May 2005. [Online]. Available: http://www.nature.com/neuro/journal/v8/n5/abs/nn1444.html

[29] F. De Martino, G. Valente, N. Staeren, J. Ashburner, R. Goebel, and E. Formisano, "Combining multivariate voxel selection and support

vector machines for mapping and classification of fMRI spatial patterns," *Neuroimage*, vol. 43, no. 1, pp. 44–58, 2008. [Online]. Available: http://dx.doi.org/10.1016/j.neuroimage.2008.06.037

[30] E. Formisano, F. De Martino, M. Bonte, and R. Goebel, ""who" is saying "what"? brain-based decoding of human voice and speech." *Science*, vol. 5903, no. 322, pp. 970–3, 2008.

[31] N. Kriegeskorte, R. Goebel, and P. Bandettini, "Information-based functional brain mapping," *Proceedings of the National Academy of Science*, vol. 103, pp. 3863–3868, 2006.

[32] N. Kriegeskorte, E. Formisano, B. Sorger, and R. Goebel, "Individual faces elicit distinct response patterns in human anterior temporal cortex," *Proceedings of the National Academy of Sciences*, vol. 104, no. 51, pp. 20 600–20 605, December 2007. [Online]. Available: http://dx.doi.org/10.1073/pnas.0705654104

[33] N. Kriegeskorte and P. Bandettini, "Combining the tools: activation- and information-based fmri analysis," *NeuroImage*, vol. 38, pp. 666–668, 2007.

[34] J.-D. Haynes, K. Sakai, G. Rees, S. Gilbert, C. Frith, and R. E. Passingham, "Reading hidden intentions in the human brain," *Current Biology*, vol. 17, no. 4, pp. 323–328, February 2007. [Online]. Available: http://dx.doi.org/10.1016/j.cub.2006.11.072

[35] S. Bode and J. Haynes, "Decoding sequential stages of task preparation in the human brain," *NeuroImage*, vol. 45, no. 2, pp. 606–613, April 2009. [Online]. Available: http://dx.doi.org/10.1016/j.neuroimage.2008.11.031

[36] J. Clithero, R. Carter, and S. Huettel, "Local pattern classification differentiates processes of economic valuation," *NeuroImage*, vol. 45, no. 4, pp. 1329–1338, 2008.

[37] M. Stokes, R. Thompson, R. Cusack, and J. Duncan, "Top-down activation of shape-specific population codes in visual cortex during mental imagery." *Journal of Neuroscience*, vol. 29, no. 5, pp. 1565–72, 2009.

[38] M. Björnsdotter and J. Wessberg, "A memetic algorithm for selection of 3d clustered features with applications in neuroscience," 2010.

[39] ——, "Particle swarm voxel clustering for multivariate fMRI mapping," *Organization for Human Brain Mapping Annual Meeting, San Fransisco, USA*, 2009.

[40] D. P. Prata, A. Mechelli, C. H. Y. Fu, M. Picchioni, T. Toulopoulou, E. Bramon, M. Walshe, R. M. Murray, D. A. Collier, and P. McGuire, "Epistasis between the dat 3' utr vntr and the comt val158met snp on cortical function in healthy subjects and patients with schizophrenia." *Proc Natl Acad Sci U S A*, vol. 106, no. 32, pp. 13 600–13 605, Aug 2009. [Online]. Available: http://dx.doi.org/10.1073/pnas.0903007106

[41] F. Benedetti, D. Radaelli, A. Bernasconi, S. Dallaspezia, A. Falini, G. Scotti, C. Lorenzi, C. Colombo, and E. Smeraldi, "Clock genes beyond the clock: Clock genotype biases neural correlates of moral valence decision in depressed patients." *Genes Brain Behav*, vol. 7, no. 1, pp. 20–25, Feb 2008. [Online]. Available: http://dx.doi.org/10.1111/j.1601-183X.2007.00312.x

[42] R. Hashimoto, T. Numakawa, T. Ohnishi, E. Kumamaru, Y. Yagasaki, T. Ishimoto, T. Mori, K. Nemoto, N. Adachi, A. Izumi, S. Chiba, H. Noguchi, T. Suzuki, N. Iwata, N. Ozaki, T. Taguchi, A. Kamiya, A. Kosuga, M. Tatsumi, K. Kamijima, D. R. Weinberger, A. Sawa, and H. Kunugi, "Impact of the disc1 ser704cys polymorphism on risk for major depression, brain morphology and erk signaling." *Hum Mol Genet*, vol. 15, no. 20, pp. 3024–3033, Oct 2006. [Online]. Available: http://dx.doi.org/10.1093/hmg/ddl244

[43] A. Krug, V. Markov, S. Krach, A. Jansen, K. Zerres, T. Eggermann, T. Stcker, N. J. Shah, M. M. Nthen, A. Georgi, J. Strohmaier, M. Rietschel, and T. Kircher, "Genetic variation in g72 correlates with brain activation in the right middle temporal gyrus in a verbal fluency task in healthy individuals." *Hum Brain Mapp*, vol. 32, no. 1, pp. 118–126, Jan 2011. [Online]. Available: http://dx.doi.org/10.1002/hbm.21005