NeuroImage xxx (2010) xxx–xxx

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Technical Note

Functional connectivity analysis of fMRI data using parameterized regions-of-interest

Wouter D. Weeda *, Lourens J. Waldorp, Raoul P.P.P. Grasman, Simon van Gaal, Hilde M. Huizenga

University of Amsterdam, Department of Psychology, The Netherlands

ARTICLE INFO

Article history:
Received 12 February 2010
Revised 7 July 2010
Accepted 9 July 2010
Available online xxxx

ABSTRACT

Connectivity analysis of fMRI data requires correct specification of regions-of-interest (ROIs). Selection of ROIs based on outcomes of a GLM analysis may be hindered by conservativeness of the multiple comparison correction, while selection based on brain anatomy may be biased due to inconsistent structure-to-function mapping. To alleviate these problems we propose a method to define functional ROIs without the need for a stringent multiple comparison correction. We extend a flexible framework for fMRI analysis (Activated Region Fitting, Weeda et al. 2009) to connectivity analysis of fMRI data. This method describes an entire fMRI data volume by regions of activation defined by a limited number of parameters. Therefore a less stringent multiple comparison procedure is required. The regions of activation from this analysis can be directly used to estimate functional connectivity. Simulations show that Activated Region Fitting can recover the connectivity of brain regions. An application to real data of a Go/No-Go experiment highlights the advantages of the method.

Introduction

Functional connectivity analysis yields insight into the brain network associated with a particular task. More specifically, it indicates which brain regions covary during task performance. An adequate functional connectivity analysis requires that all activated brain regions are incorporated; otherwise the analysis will yield a biased indication of the brain network associated with a task (Eichler, 2005).

Regions of task-related activation are commonly identified in two ways. First, they may be derived from a general linear model (GLM) analysis, where above threshold voxels are used as regions-of-interest (ROIs). Identifying regions this way may miss important regions as most GLM multiple testing corrections tend to be conservative (Nichols and Hayasaka, 2003). This is especially problematic in situations where there is a low signal-to-noise ratio (SNR), for example in single-subject or developmental studies. This problem of conservativeness may be alleviated by including a localizer scan. Localizer scans elicit more significant activation; however, the use of functional localizers is still subject to debate (for example, Friston et al. (2006)). Second, regions of task-related activity can be formulated a-priori. This however requires that the constituents of the brain network are already known, which is often not the case. Moreover, it requires between-subjects homogeneity of function-to-structure mapping, which is often hard to attain. That is, functional regions tend to map to different structural locations across subjects (Hunton et al., 1996).

In this paper we develop a method that has increased sensitivity to identify active brain regions, which in turn yields a more precise indication of the task-related network. The method is based on a generalization of Activated Region Fitting (ARF, Weeda et al., 2009). In ARF the amplitudes of an unthresholded GLM analysis are described by a parsimonious spatial model in which activated regions are parameterized by their location, spatial extent and amplitude. The number of parameters in this model is therefore low compared to the number of voxels. Consequently ARF has increased power to detect activation (Weeda et al., 2009). The ARF framework can be easily extended from a localization perspective to a connectivity perspective. By regressing trial specific fMRI data on activated regions, the trial-by-trial varying amplitude of activated regions can be estimated. These trial-by-trial amplitudes can be used to estimate functional connectivity between activated brain regions (cf. Rissman et al., 2004; Fox et al., 2006). The main advantage of the ARF framework is that it yields a more realistic description of the constituents of a functional brain network, since it is not required to specify regions-of-interest a-priori or to rely on a GLM analysis that is stringently corrected for multiple comparisons.

The organization of the paper is as follows. First we describe the details of the ARF method and its extension to connectivity analysis. Second, we test the ARF connectivity method in simulations. Third, we illustrate the method by applying it to single-subject data obtained in an unconscious Go/No-Go study (cf. van Gaal et al., 2010). Finally, we discuss limitations and possible extensions.

Methods

The data used in the ARF framework are *beta*- or *t*-values of a single-subject or whole-group GLM analysis. These *beta*- or *t*-values are described by a spatial model that consists of multiple activated regions. The number of required regions is determined by the Bayesian

* Corresponding author. University of Amsterdam, Department of Psychology, Roetersstraat 15, 1018 WB Amsterdam, The Netherlands.
E-mail address: w.d.weeda1@uva.nl (W.D. Weeda).

1053-8119/$ – see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2010.07.022

Information Criterion (BIC). The BIC penalizes the fit for the number of parameters, keeping the number of regions in the spatial model low. Hypothesis tests are performed on the parameters of each activated region, allowing for hypotheses on location, spatial extent, and amplitude. A full ARF analysis fits models with different numbers of regions and selects the model with the smallest BIC value. Thereafter hypothesis test on the amplitude and spatial extent parameters of the selected model are performed. Regions with an amplitude and/or extent that do not deviate from zero are omitted from the final model. Details of the procedure can be found in Weeda et al. (2009)[1].

In Weeda et al. (2009) the method was restricted to two dimensions; that is single-slice or flattened maps. The current implementation extends the ARF method to three-dimensional full volume analysis. For purposes of clarity the details of the 3D extension are described in detail in Appendix A.

Connectivity analysis

Within the ARF framework activity is modeled by multiple activated regions. The framework is well suited for connectivity analysis, as the identified regions can be directly used as regions-of-interest (ROIs). The connectivity analysis consists of two steps. First, the single-trial data are regressed on estimated activated regions obtained from the ARF analysis. This creates trial-by-trial amplitude estimates for each activated region. Second, these trial-by-trial amplitudes are used to calculate functional connectivity estimates between regions. Details of the two steps are explained below.

Single-trial amplitude estimation

Single-trial *data* are obtained by creating a model with a regressor for each single-trial, convolved with a Hemodynamic Response Function (HRF, cf. Boynton et al., 1996). The time-series are regressed on this model using a standard GLM analysis. The *beta* values from this analysis constitute the single-trial data (cf. Rissman et al., 2004), which are input to the connectivity analysis.[2]

Let $\mathbf{f}_j$ be the $(N\times 1)$ vector of ARF model estimates of region $j=1,\ldots,J$, where J denotes the number of regions and N denotes the number of voxels. We then construct an $(N\times J)$ design matrix F with in each column the model estimates $\mathbf{f}_j$ for region j with unit amplitude. Note that no spatial boundary for the estimated regions is set because the contribution to the model of voxels far from the center is nearly zero. The design matrix F is used to obtain single-trial estimates of the amplitude of the activated regions.

Let $\mathbf{y}_k$ be an $(N\times 1)$ vector containing the single-trial data for all voxels of trial $k=1,\ldots,K$. For each trial k the data $\mathbf{y}_k$ are regressed on the model F using linear regression:

$$\mathbf{y}_k = F\gamma_k + \varepsilon_k. \tag{1}$$

The vector of estimated single-trial amplitudes $\hat{\gamma}_k$ of J regions is then derived from:

$$\hat{\gamma}_k = (F'F)^{-1}F'\mathbf{y}_k. \tag{2}$$

Then the $(J\times K)$ matrix $\mathbf{G}$ is constructed such that it has in each column the single-trial amplitudes $\hat{\gamma}_k$.

Connectivity estimation

Functional connectivity estimates are obtained by correlating the rows of $\mathbf{G}$. The resulting $(J\times J)$ matrix $\mathbf{M}$ contains the correlations between regions. Given that the ARF parameter estimates are close to their optimal values, the correlation estimates in $\mathbf{M}$ are consistent (see Appendix C).

Testing differences in connectivity

Differences in connectivity between conditions can be tested by converting correlations to z-scores, calculating the difference between these z-scores, and testing whether this difference is significant (Myers and Well, 2003; ch. 18.3). The Fischer-z transformation is used to convert each element $m_{r,c}$ of $\mathbf{M}$ to z-scores (Fisher, 1915):

$$z_{r,c} = \frac{1}{2}\ln\left[\frac{1+m_{r,c}}{1-m_{r,c}}\right]. \tag{3}$$

The difference of z-scores of conditions A and B is then calculated using:

$$z_{diff} = \frac{z^A_{r,c} - z^B_{r,c}}{\sqrt{\frac{1}{K^A-3}+\frac{1}{K^B-3}}} \tag{4}$$

which, under the null hypothesis of no differences, is standard normally distributed. This tests the null hypothesis that the connectivity between regions r and c is equal across conditions A and B.

Simulations

In simulations we investigate whether the ARF connectivity framework can recover the correlation structure in the data. In the simulations we varied several factors that may affect performance of ARF connectivity estimates. We varied the size of regions, the shape of regions and the amount of spatial smoothing applied to the data. In addition we varied the amount of single-trial variation in the location of regions and in the extent of regions.

In all conditions we created 3D fMRI volumes with three activated regions. All regions were placed far apart to ensure a minimum of spatial overlap. White noise was added according to the method described in Weeda et al. (2009, Appendix A), using different levels (.1, .5, 1, and 2) of single-trial signal-to-noise ratios (SNRs) commonly found in fMRI studies (Huettel et al., 2001). For each trial the amplitude of the signals was sampled from a multivariate normal distribution with correlation matrix[3]:

$$\begin{bmatrix} 1 & .5 & .7 \\ .5 & 1 & .35 \\ .7 & .35 & 1 \end{bmatrix}$$

to simulate a range of correlations between regions. The simulations were performed 100 times for each SNR level. The variance–covariance matrix defining the simulated regions was:

$$\begin{bmatrix} \theta_x^2 & .01\cdot\theta_x\theta_y & -.1\cdot\theta_x\theta_z \\ .01\cdot\theta_x\theta_y & \theta_y^2 & .1\cdot\theta_y\theta_z \\ -.1\cdot\theta_x\theta_z & .1\cdot\theta_y\theta_z & \theta_z^2 \end{bmatrix}$$

Details of the simulation settings in each condition are given below.

Size

The effect of size was investigated by simulating three Gaussian shaped regions with different θ_x, θ_y, and θ_z parameters. The simulated regions either contained 721 voxels within the 95% isocontour

[1] The ARF procedure is available as an open-source package for R (R Development Core Team, 2009) and can be downloaded from the authors' website: http://home.medewerker.uva.nl/w.d.weeda1.

[2] To compare performance of this regression approach with an approach using raw time-series, simulations were performed. Details can be found in Appendix B.

[3] The variance of the amplitude parameters over trials was set to 16,000 as estimated from the real data by van Gaal et al. (2010).

Please cite this article as: Weeda, W.D., et al., Functional connectivity analysis of fMRI data using parameterized regions-of-interest, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.07.022

(denoted as size = 2) or contained 2477 voxels within the 95% isocontour (denoted as size = 3).

Extent variation

Extent variation was simulated for size = 2 by introducing single-trial variation in the extent parameters θ_x, θ_y, and θ_z, which were sampled from a normal distribution with sd of .25 or .50. This created regions with varying extent ranging between 309 and 1389 voxels in the sd = .25 condition and between 85 and 2477 voxels in the sd = .5 condition.

Location variation

Location variation was simulated for size = 2 by introducing single-trial variation in the center of regions, this variation was sampled from a normal distribution with sd = 1 or sd = 2 for *x*, *y* and *z* directions. This created location variation between −2 and 2 voxels when sd = 1 or between −4 and 4 voxels when sd = 2.

Smoothing

The effect of smoothing was investigated for size = 2. The single-trial data was smoothed by an FWHM filter of 2 times voxel size or an FWHM filter of 3 times voxel size.

Shape

The effect of region shape was investigated by taking three anatomical shapes from the HarvardOxford Probabilistic Atlas (from the FSL software package, Smith et al., 2004). The regions were the left superior temporal gyrus (STG, 1250 voxels), the left inferior frontal gyrus (IFG, 2918 voxels) and the left lateral occipital cortex (LOC, 4318 voxels). See Fig. 2, bottom-right panel, for an example of these regions. The correlations between regions were all set to .5.

ARF connectivity estimation and standard methods

In each simulation, first an ARF model was fitted to averaged data. In the simulations it was correctly assumed that three regions were active. Then, using the model predictions, correlations between regions were estimated as described in the Methods section. We compared the connectivity estimates obtained from ARF with standard methods: connectivity estimates derived from average activity in ROIs and connectivity estimates based on an eigenvariate approach (Friston et al., 2006). In the latter two approaches each ROI consisted of voxels within the 95% isocontour of the known simulated region.

Note that for the ARF simulations it is assumed that the number of regions is known precisely. This yields an upper bound of ARF performance. In the 'average ROI' analysis both the number and the location of regions were known. Therefore, the average ROI analysis also provides an upper bound of performance. In the 'eigenvariate' analysis both the number and location of ROIs were again assumed to be known exactly. Also, the first eigenvariate was assumed to be known exactly (i.e. estimated without error), thus providing an upper bound of the precision that can be obtained. Note in addition that the comparisons favor the 'average ROI' and 'eigenvariate' analyses over the ARF method since in the standard methods both number and location of regions are assumed to be known (and even the first eigenvariate in the 'eigenvariate' approach), whereas in the ARF method only the number of regions is assumed to be known.

Power

In the aforementioned simulations the number (ARF) or the number and location (standard methods) of regions was known precisely. In order to compare the power of these approaches (see also Weeda et al. (2009)) we performed a final set of simulations in which the full ARF procedure was compared to False Discovery Rate (FDR, Genovese et al., 2002) thresholding. In these simulations we reanalyzed the dataset of the size = 3 simulations using a full ARF procedure: the method selects the number of regions according to the lowest BIC value and subsequently tests amplitude and extent of these regions using the Wald test (with $p<.05$). The ARF method was contrasted with standard ROI detection using FDR correction, using $p<.05$ and a minimum of 10 voxels above threshold within a $(5\times5\times5)$ box around the center of the simulated region.

Results

Fig. 1 shows results for simulations in which size, location variation, extent variation, smoothing and signal-to-noise ratio were varied[4]. Results show that correlations are severely attenuated by noise in the average ROI analysis. This effect is strongest for the higher correlations (see Fig. 1 and Supplementary Figs. 1 and 2). The ARF method outperforms the average ROI method in all cases, especially in situations where the signal-to-noise ratio is low. In the size simulations (Fig. 1a) both methods perform better when region size is increased; there is no indication that size differentially affects both methods. In the location variation simulations (Fig. 1b) performance of both methods decreases with increased location variation. ARF performs better than the average ROI method under small location variation; this advantage decreases with increased location variation. In the varying extent simulations (Fig. 1c) performance of both methods decreases with increased extent variation. The advantage of the ARF method is especially pronounced when extent variation is large. The smoothing simulations (compare Fig. 1a, left hand panel (unsmoothed data) to Fig. 1d (smoothed data)) indicate that the advantage of the ARF method decreases with increased smoothing.

Fig. 2 shows the results for simulations with realistically shaped regions. It can be seen that the ARF method outperforms especially the standard averaging method and to a lesser extent the eigenvariate method.

Power

Results for the full ARF procedure are in Fig. 3. As can be seen the ARF procedure outperforms the FDR method in the lower SNR conditions (SNR = .1 and SNR = .5). In the SNR = .1 condition ARF selects the correct model in 21% of cases, while FDR detects no regions at all. For the SNR = .5 condition ARF selects the correct model in all cases, while FDR only detects all regions in 41% of cases. For SNR = 1 and SNR = 2 both methods detect all regions.

False positive rate

To assess the false positive rate (i.e. the chance of detecting an active region when there is no activity at all) of the ARF method we simulated 100 null datasets that contained only noise. We fitted an ARF model with three regions and tested with the Wald statistic (cf. Appendix A) how many regions had significant ($p<.05$) amplitude and extent parameters. The false positive rate was as required below 5%; that is 3%.

Empirical application

In order to illustrate the ARF method we analyzed a single subject out of a dataset from a Go/No-Go experiment performed by van Gaal et al. (2010). The subject performed a Go/No-Go task in which

[4] Only results from the correlations of .70 are shown. Effects for the correlations of .35 and .5 are comparable and can be found in Figs. 1 and 2 of the Supplementary Materials. Also, the ARF approach and the eigenvariate method produced comparable results. Eigenvariate results were therefore omitted from the plots in Fig. 1.

Please cite this article as: Weeda, W.D., et al., Functional connectivity analysis of fMRI data using parameterized regions-of-interest, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.07.022

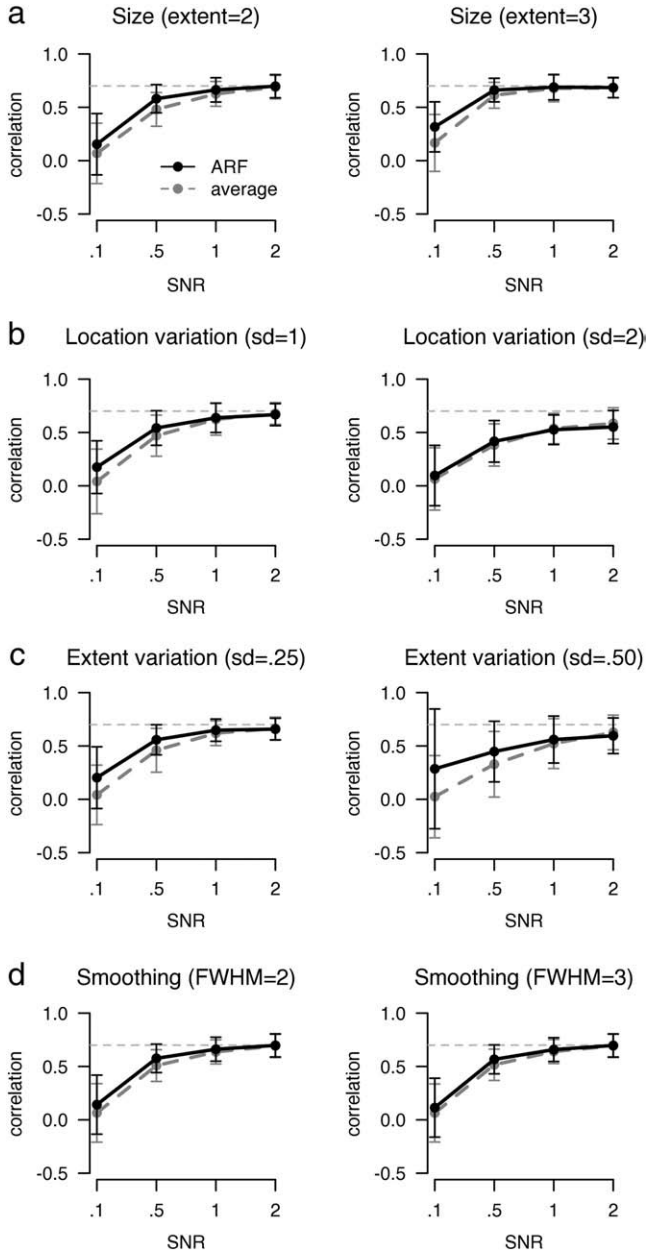

Fig. 1. Functional connectivity estimates of the simulated data as a function of SNR. Solid lines indicate estimates for the ARF method. Dashed lines indicate estimates based on the average of an ROI. Error-bars indicate the 95% confidence interval of the estimates. Panel (a) shows size variation, (b) shows location variation, (c) shows extent variation and (d) shows smoothing variation.

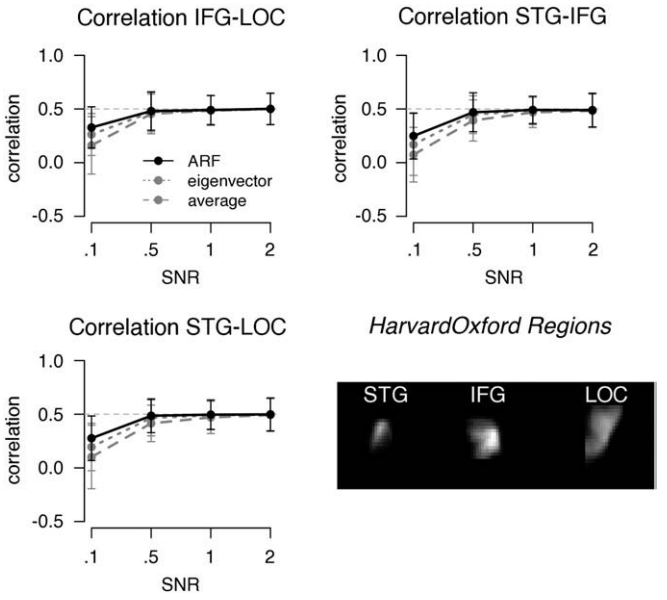

Fig. 2. Functional connectivity estimates of the realistic shape simulations as a function of SNR. Solid lines indicate estimates for the ARF method. Dashed lines indicate estimates based on the average of an ROI, dotted lines indicate estimates based on the first eigenvariate. Examples of the regions are depicted in the lower-right panel.

conscious No-Go trials (weakly masked No-Go cue) and unconscious No-Go trials (strongly masked No-Go cue) were mixed randomly with Go trials. Pre-processing followed van Gaal et al. with one exception that the data were not spatially smoothed. A contrast of two conditions was analyzed: the strongly masked Go condition where the subject responded (Go) versus the strongly masked No-Go condition where the subject responded (No-Go). The contrast of No-Go>Go highlighted the 'unconscious No-Go network' (cf. van Gaal et al.). This contrast was used to estimate the activated regions with ARF. After demeaning the raw time-series to remove global effects, the single-trial responses to each stimulus were estimated by regressing the time-series to the convolved response function (double-gamma HRF; Glover, 1999) of each trial. These data were used to estimate the functional connectivity in the No-Go and Go conditions with the ARF estimates.

Results

The No-Go>Go contrast was not significant when thresholded at $p = .05$ using Bonferroni correction. With False Discovery Rate (FDR) correction (Genovese et al., 2002), only three (adjacent) voxels in the left frontal cortex were significant[5]. Running an ARF analysis revealed an optimal model of 23 activated regions of which 22 regions had significant amplitude and spatial extent parameters (Bonferroni corrected). Mostly, activated regions were located in the parietal areas, motor areas (including the (pre) supplementary motor areas and motor cortex) and frontal areas (including the inferior frontal gyrus), consistent with typical results reported in group Go/No-Go studies (cf. van Gaal et al., 2010). Further, replicating several neuroimaging studies, activated regions were found also in some sub-cortical areas (including cerebellum and globus pallidus) thought to be involved in motor control and response inhibition (Aron et al., 2007; Chambers et al., 2009). Parameter estimates of the regions are reported in Table 1 of the Supplementary Materials (note that these regions correspond roughly to the regions reported previously on a group level; see Supplementary Table 2 of van Gaal et al. (2010)).

The analysis of differences in functional connectivity strength between No-Go and Go conditions revealed four (commonly observed) regions where correlations were significantly higher in the No-Go condition than in the Go condition. Fig. 4 shows the location of these four regions; Fig. 5 shows functional connectivity estimates of these four regions. First of all the correlation between the pre-supplementary motor area (pSMA) and motor cortex (MC) increased from −.03 (Go) to −.35 (No-Go). Further, the correlation between pSMA and left cerebellum (L-Cb) increased from −.34 (Go) to −.58 (No-Go), whereas the correlation between the globus pallidus (GP; output structure of the basal ganglia) and right cerebellum (R-Cb) increased from .58 (Go) to .77 (No-Go).

Although the exact nature and functional interpretation of these increased connectivity measures is not the key issue of this paper,

[5] The Bonferroni and FDR thresholding was also applied to smoothed data (FWHM = 5 mm), leading to approximately the same results.

Please cite this article as: Weeda, W.D., et al., Functional connectivity analysis of fMRI data using parameterized regions-of-interest, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.07.022

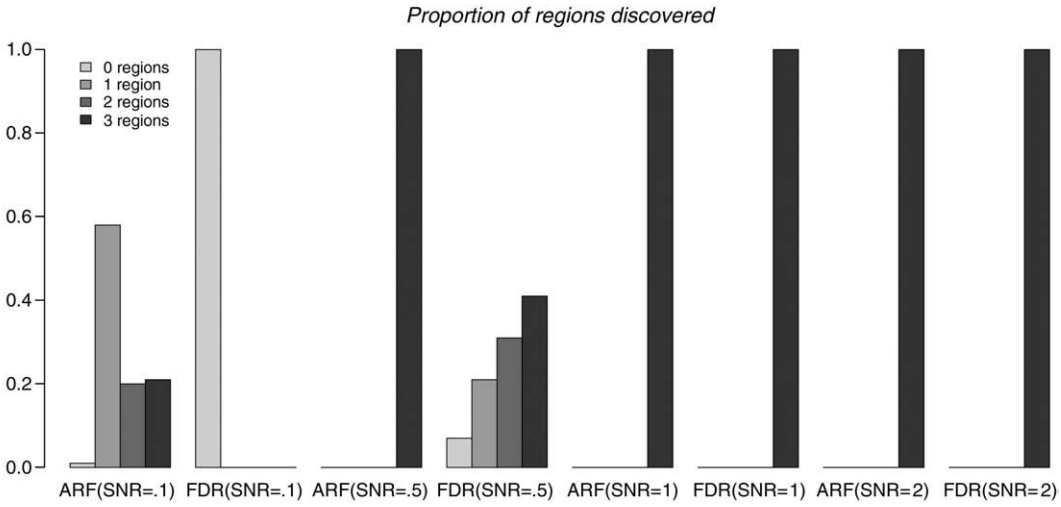

Fig. 3. Proportion of regions detected by the ARF and FDR methods. Bars indicate the proportion of regions detected by both methods as a function of SNR (lightest bar indicates 0 regions detected; darkest bar indicates 3 regions detected).

previous research has identified a crucial role of the pSMA and basal ganglia (including the GP) in response inhibition (Aron et al., 2007). In addition, various studies have reported involvement of the Cb in motor control as well as response inhibition on No-Go trials (Chambers et al., 2009).

Discussion

Functional connectivity analysis with ARF can be a helpful tool in the analysis of interactions between brain regions. Its most prominent advantage is that a more realistic indication of regions in a network can be obtained, since it is not required to specify regions-of-interest a-priori or to rely on a GLM analysis that is stringently corrected for multiple comparisons. Application of the full ARF procedure shows that the method has increased sensitivity to detect ROIs in low SNR conditions. In addition, results of simulations show that the method can recover correlations between brain regions. Although estimates are attenuated in low SNR conditions, ARF connectivity estimation outperforms the average ROI method in all cases. Even in conditions with spatially misspecified models (anatomical shape condition), ARF outperforms the average ROI and eigenvariate methods.

The use of spatial models in the ARF method makes it less sensitive to noise than averaging activity within an ROI. More specifically, as the ARF method weights activation by the spatial model, noise at edges of ROIs does not influence estimation as much as in standard averaging. In cases with anatomically shaped regions ARF also outperforms an eigenvariate approach. This difference may be due to the difference in weighting between the methods. The eigenvariate approach assigns weights to voxels that are a function of the amount of temporal variance explained by these voxels; these weights are independent of physical location. ARF weights activation dependent on the physical location: activity in the center is weighted highest and weights decrease as the distance from the center increases.

Applying the ARF method to single-subject data revealed activation in areas consistent with many previous Go/No-Go studies (cf. van Gaal et al., 2010) while standard thresholding in this single subject yielded only one very small area of activation. Connectivity estimates showed potentially insightful differences in connectivity strength between regions typically associated with response inhibition when comparing Go and No-Go conditions.

The ARF method uses an approach similar to that of Rissman et al. (2004) to calculate trial-by-trial data. Estimation of trial-by-trial

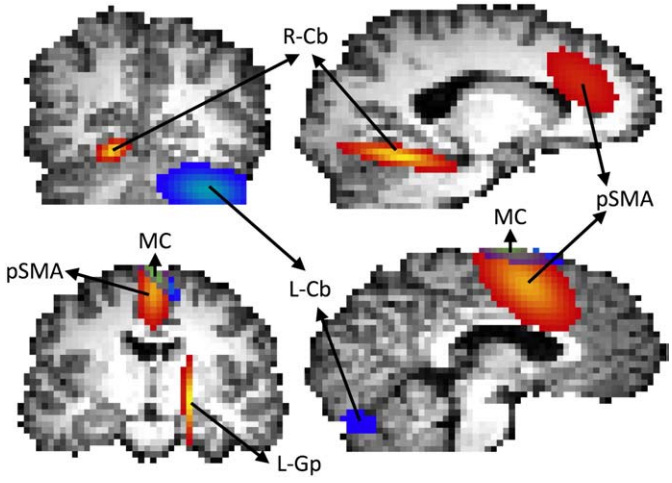

Fig. 4. Location of the activated regions (pre-) supplementary motor area (pSMA), motor cortex (MC), left (L) and right (R) cerebellum (Cb), and globus pallidus (Gp). Regions in red-yellow indicate activity where No-Go>Go. Regions in blue-green indicate activity where Go>No-Go. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

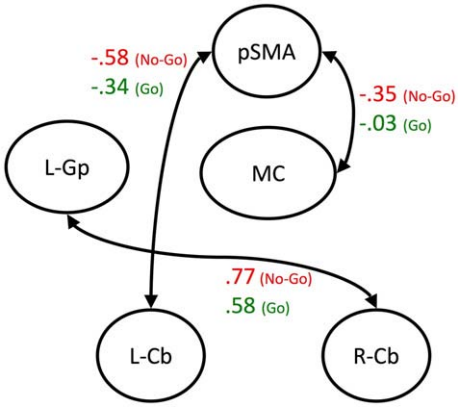

Fig. 5. Functional connectivity with absolute connectivity strength greater in the No-Go condition than the Go condition ($p<.05$). Numbers in red indicate correlation in the No-Go condition; numbers in green indicate correlation in the Go condition. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Please cite this article as: Weeda, W.D., et al., Functional connectivity analysis of fMRI data using parameterized regions-of-interest, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.07.022

amplitudes is different as ARF uses spatial models to estimate these amplitudes. Also, in comparison with the methods of Rissman et al. and Rissman et al. (2004), ARF requires no a-priori specification of ROIs.

This study leaves room for further exploration. First, correlations are underestimated in low SNR conditions (cf. Spearman, 1904). This attenuation of correlations due to noise is reduced in the ARF analysis, but is still present. On the one hand this finding is comforting. Since the estimates are conservative the chances of finding spurious correlations are reduced. On the other hand, weak functional couplings between activated brain areas might be missed. Future work may consider extended ARF based methods that account for attenuation due to noise. These methods should explicitly model the noise (co)variance structure.

Second, the ARF framework is not limited to functional connectivity estimation. The method can also be used to identify ROIs for effective connectivity analyses like Dynamic Causal Modeling (Friston et al., 2003), Structural Equation Modeling (McIntosh and Gonzalez-Lima, 1994; Gonçalves and Hall, 2003), or graphical modeling (Ramsey et al., 2010). The ARF framework thus provides a powerful, flexible framework for the localization of — and the analysis of functional connectivity between — activated brain areas.

Acknowledgments

The authors wish to thank two anonymous reviewers for their comments. This research was funded by the Netherlands' Organization for Scientific Research (NWO).

Appendix A. Extension to full volume fMRI analysis

The method of Activated Region Fitting (ARF) was first developed for flattened activity maps of fMRI data (Weeda et al., 2009). In this paper the method is extended to full volume fMRI analysis. This extension encompasses changes to the spatial model and the test statistics. In the paragraphs below the details of the changes are explained.

Spatial model

Activated Region Fitting uses a spatial model with multiple activated regions to describe an entire volume of fMRI data. A region is modeled by a multivariate Gaussian. The three-dimensional model for voxel n with $n=1,\ldots,N$ voxels and $j=1,\ldots J$ regions is:

$$f(\mathbf{x}_n,\theta)=\sum_{j=1}^{J}\frac{\theta_{10j}}{(2\pi)^{3/2}\left|\Sigma_j\right|^{1/2}}\exp\left[-\frac{1}{2}\left(\mathbf{x}_n-\mathbf{k}_j\right)'\Sigma_j^{-1}\left(\mathbf{x}_n-\mathbf{k}_j\right)\right]. \quad (5)$$

The location of voxel n is contained in vector $\mathbf{x_n}=(x_n,y_n,z_n)'$. The parameters for the location of the center of region j are in vector $\mathbf{k_j}=(\theta_{1j},\theta_{2j},\theta_{3j})'$. The parameters defining the shape of region j are in matrix Σ_j:

$$\Sigma_j=\begin{bmatrix}\theta_{4j}^2 & \theta_{4j}\theta_{5j}\theta_{7j} & \theta_{4j}\theta_{6j}\theta_{8j}\\ \theta_{4j}\theta_{5j}\theta_{7j} & \theta_{5j}^2 & \theta_{5j}\theta_{6j}\theta_{9j}\\ \theta_{4j}\theta_{6j}\theta_{8j} & \theta_{5j}\theta_{6j}\theta_{9j} & \theta_{6j}^2\end{bmatrix}.$$

In Eq. (5) $|\Sigma_j|$, the determinant of Σ_j, is the spatial extent (i.e. volume) of region j. Finally, the amplitude of region j is defined by θ_{10j}. Parameter estimation, model selection and estimation of the covariance matrix of the parameter estimates is performed as in Weeda et al. (2009). Parameters are estimated using Weighted Least Squares. The number of regions is determined by using the Bayesian Information Criterion (BIC). To accommodate model misspecification (in the number of spatial models to describe a volume as well as misspecification in the shape of each region) a robust estimator of the parameter covariance matrix is obtained from the Sandwich estimator (Weeda et al., 2009; Waldorp, 2009).

Hypothesis testing

Let $\mathbf{a}\left(\hat{\theta}_j\right)$ contain the hypotheses of interest, for example the 5 null hypotheses: x location is equal to c_1, y location is equal to c_2, z location is equal to c_3, spatial extent of the region is zero, and amplitude of the region is zero:

$$\mathbf{a}'(\hat{\theta}_j)=\begin{bmatrix}\theta_{1j}-c_1 & \theta_{2j}-c_2 & \theta_{3j}-c_3 & |\sum_j| & \theta_{10j}\\ x\_location & y\_location & z\_location & extent & amplitude\end{bmatrix}$$

Let matrix $\hat{\mathbf{A}}_j$ contain the associated first-order derivatives to the parameters in our example:

$$\hat{\mathbf{A}}_j=\begin{bmatrix}1&0&0&0&0&0&0&0&0&0\\0&1&0&0&0&0&0&0&0&0\\0&0&1&0&0&0&0&0&0&0\\0&0&0&\frac{\partial|\Sigma_j|}{\partial\theta_{4j}}&\frac{\partial|\Sigma_j|}{\partial\theta_{5j}}&\frac{\partial|\Sigma_j|}{\partial\theta_{6j}}&\frac{\partial|\Sigma_j|}{\partial\theta_{7j}}&\frac{\partial|\Sigma_j|}{\partial\theta_{8j}}&\frac{\partial|\Sigma_j|}{\partial\theta_{9j}}&0\\0&0&0&0&0&0&0&0&0&1\end{bmatrix}.$$

Let matrix $\mathbf{C}_j$ denote the covariance matrix of the parameter estimates for region j. The test statistic then equals:

$$\mathbf{a}'(\hat{\theta}_j)\left[\hat{\mathbf{A}}_j\mathbf{C}_j\hat{\mathbf{A}}_j'\right]\mathbf{a}(\hat{\theta}_j). \quad (6)$$

This test statistic is under the null hypothesis asymptotically F distributed with N and $(N-p)$ degrees of freedom (with p being the number of parameters). This yields a multivariate test; univariate tests are obtained by using the appropriate elements of $\mathbf{a}\left(\hat{\theta}_j\right)$ and $\hat{\mathbf{A}}_j$ separately in Eq. (6).

Appendix B. Raw time-series versus trial-by-trial amplitude

Single-trial *beta* values constitute the input data for the ARF procedure. An alternative way is to use raw time-series as input. On request of one of the reviewers, we compare these two approaches in a small simulation study.

For each voxel in a Gaussian shaped ROI we simulated raw time-series of 2400 s containing 48 non-overlapping BOLD responses. The raw time-series were created by convolving a stick function with a standard HRF. The weights of the stick function differed to create a time-series containing BOLD responses with different amplitudes. These raw time-series were weighted by the (spatial) Gaussian function and white noise was added under four SNR conditions (.1, .5, 1 and 2). For 100 datasets we simulated two ROIs in which the simulated amplitudes of the BOLD responses (that is, the weights of the stick function) of the ROIs correlated .5. This led to a correlation of the raw time-series of around .92. The higher correlation is due to the BOLD response being relatively slow and therefore highly correlated in time.

To each dataset we fitted an ARF model with two regions. Using the ARF model estimates we calculated connectivity using the ARF method as described in the Methods section. We contrasted the ARF method with an eigenvariate approach (Friston et al., 2006) (i) using raw time-series, and (ii) using the single-trial *beta* values (Rissman et al., 2004). The ROIs used in this approach were calculated using the voxels in the 95% confidence interval of the regions estimated by ARF.

Results are shown in Fig. 6. As can be seen the ARF method outperforms the eigenvariate approach using the *beta* values in the SNR = .1 and SNR = .5 conditions. Fig. 6 also shows that estimates

Please cite this article as: Weeda, W.D., et al., Functional connectivity analysis of fMRI data using parameterized regions-of-interest, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.07.022

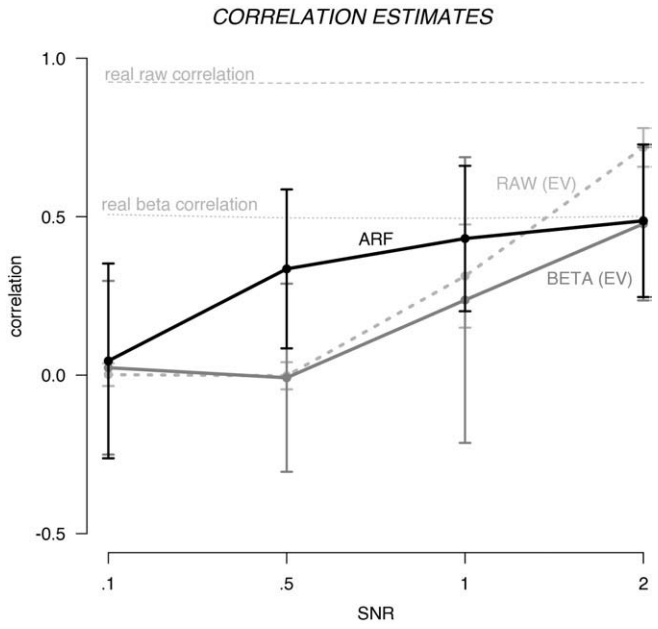

Fig. 6. Connectivity estimates of the simulated data as a function of SNR. The solid black line indicate estimates for the ARF method. Solid gray line indicates estimates calculated using the eigenvariate approach on *beta* values. Dashed gray line indicates estimates calculated using the eigenvariate approach on raw time-series. Error-bars indicate the 95% confidence interval of the estimates.

based on the raw time-series are attenuated by noise in all conditions; that is, they never equal the true value of .92.

Appendix C. Consistency of correlation estimates

We need to show that the estimates of the correlations between regions are consistent; that is, that the estimated correlation between regions approaches the optimal correlation as the number of trials tends to infinity. Our assumptions are (i) that the least squares (LS) estimate of the ARF region parameter $\hat{\theta}_n$ is consistent for the optimal value θ_* which minimizes the expectation of the LS function $Q(\theta)$

$$E\{Q(\theta)\} = \int (y - f(\theta))' W^{-1} (y - f(\theta)) g(y) dy$$

where y contains the averaged data (that is the *beta* or *t*-values from a GLM analysis), W contains the variances of the *beta* or *t*-values and $f(\theta)$ contains the Gaussian model with parameters θ; and (ii) that y has finite variance $n^{-1/2}\Psi$. Note that we do not assume that the Gaussian ARF model is correct (see White (1982) for examples).

The LS estimate $\hat{\theta}_n$ is then within $o_p(n^{-1/2})$ of θ_*. Since the first-order partial derivative of the Gaussian model in $F(\theta)$ is continuous (Weeda et al., 2009), by (i) we have that $F(\hat{\theta}_n)$ tends to $F(\theta_*)$ as the number of trials n goes to infinity. By assumption (ii), the estimate in Eq. (2) $\hat{\gamma}_k = (F'F)^{-1}F'\mathbf{y}_k$ for each trial k has finite variance $(F'F)^{-1}F'\Psi F(F'F)^{-1}$. Additionally, it is easily seen that the estimate $\hat{\gamma}_k$ is asymptotically normally distributed with variance $\Omega = (F'F)^{-1}F'\Psi F(F'F)^{-1}$. That means that each $\hat{\gamma}_k$ is a realization of the random variable $N(\gamma_*, \Omega)$. Therefore, the sample correlations in $\mathbf{M}$ obtained from the γ_k are consistent for the true correlations, as in any case with normally distributed variables (see e.g., Bilodeau and Brenner (1999)).

Appendix D. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.neuroimage.2010.07.022.

References

Aron, A.R., Durston, S., Eagle, D.M., Logan, G.D., Stinear, C.M., Stuphorn, V., 2007. Converging evidence for a fronto-basal-ganglia network for inhibitory control of action and cognition. J. Neurosci. 27, 11860–11864.
Bilodeau, M., Brenner, D., 1999. Theory of Multivariate Statistics. Springer-Verlag, New York, NY.
Boynton, G.M., Engel, S.A., Glover, G.H., Heeger, D.J., 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. J. Neurosci. 16, 4207–4221.
Chambers, C.D., Garavan, H., Bellgrove, M.A., 2009. Insights into the neural basis of response inhibition from cognitive and clinical neuroscience. Neurosci. Biobehav. Rev. 33, 631–646.
Eichler, M., 2005. A graphical model for evaluating effective connectivity in neural systems. Philosophical Transactions of the Royal Society B 360, 953–967.
Fisher, R.A., 1915. Distribution of the values of the correlation coefficient in samples form an indefinitely large population. Biometrika 10, 507–521.
Fox, M.D., Snyder, A.Z., Zacks, J.M., Raichle, M.E., 2006. Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses. Nat. Neurosci. 9, 23–25.
Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. Neuroimage 19, 1273–1302.
Friston, K.J., Rothstein, P., Geng, J.J., Sterzer, P., Henson, R.N., 2006. A critique of functional localizers. Neuroimage 30, 1077–1087.
Genovese, C.R., Lazar, N.A., Nichols, T., 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. Neuroimage 15, 870–878.
Glover, G.H., 1999. Deconvolution of impulse response in event-related BOLD fMRI. Neuroimage 9, 416–429.
Gonçalves, M.S., Hall, D.A., 2003. Connectivity analysis with structural equation modeling: an example of voxel selection. Neuroimage 20, 1455–1467.
Huettel, S.A., Singerman, J.D., McCarthy, G., 2001. The effects of aging upon the hemodynamic response measured by functional MRI. Neuroimage 13, 161–175.
Hunton, D.L., Miezin, F.M., Buckner, R.L., van Mier, H.I., Raichle, M.E., Petersen, S.E., 1996. An assessment of functional–anatomical variability in neuroimaging studies. Hum. Brain Mapp. 4, 122–139.
McIntosh, A.R., Gonzalez-Lima, F., 1994. Structural equation modeling and its application to network analysis in functional brain imaging. Hum. Brain Mapp. 2, 2–22.
Myers, J.L., Well, A.D., 2003. Research Design and Statistical Analysis. Lawrence Erlbaum Associates, Mahwah, NJ.
Nichols, T., Hayasaka, S., 2003. Controlling the familywise error rate in functional neuroimaging: a comparative review. Stat. Meth. Med. Res. 12, 419–446.
R Development Core Team, 2009. R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. www.R-project.org. 2009. ISBN 3-900051-07-0.
Ramsey, J.D., Hanson, S.J., Hanson, C., Halchenko, Y.O., Poldrack, R.A., Glymour, C., 2010. Six problems for causal inference from fMRI. Neuroimage 49, 1545–1558.
Rissman, J., Gazzaley, A., D'Esposito, M., 2004. Measuring functional connectivity during distinct stages of a cognitive task. Neuroimage 23, 752–763.
Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., Niazy, R., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., Matthews, P.M., 2004. Advances in functional and structural MR image analysis and implementation as FSL. Neuroimage 23, 208–219.
Spearman, C., 1904. On the proof and measurement of association between two things. Am. J. Psychol. 15, 72–101.
van Gaal, S., Ridderinkhof, K.R., Scholte, H.S., Lamme, V.A.F., 2010. Unconscious activation of the prefrontal No-Go network. J. Neurosci. 30, 4143–4150.
Waldorp, L.J., 2009. Robust and unbiased variance of GLM coefficients for misspecified autocorrelation and hemodynamic response models in fMRI. Int. J. Biomed. Imaging 2009 11 pages.
Weeda, W.D., Waldorp, L.J., Christoffels, I., Huizenga, H.M., 2009. Activated region fitting: a robust high-power method for fMRI analysis using parameterized regions of activation. Hum. Brain Mapp. 30, 2595–2605.
White, H., 1982. Maximum likelihood estimation of misspecified models. Econometrica 50, 1–25.

Please cite this article as: Weeda, W.D., et al., Functional connectivity analysis of fMRI data using parameterized regions-of-interest, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.07.022