Intensity-Based Masking: A Tool to Improve Functional Connectivity Results of Resting-State fMRI

Michael Peer,[1,2]* Sami Abboud,[1,3] Uri Hertz,[1,4] Amir Amedi,[1,5,6] and Shahar Arzy[1,2]

[1]*Department of Medical Neurobiology, the Institute for Medical Research Israel-Canada, Faculty of Medicine, Hadassah Hebrew University Medical School, Jerusalem 91120, Israel*
[2]*Department of Neurology, Hadassah Hebrew University Medical Center, Jerusalem 91120, Israel*
[3]*Sorbonne Universités, UPMC Univ Paris 06 UMR S 1127, Institut du Cerveau et de la Moelle épinière, ICM, Inserm U 1127, CNRS UMR 7225, Paris F-75013, France*
[4]*UCL Institute of Cognitive Neuroscience, University College London, Alexandra House, London WC1N 3AR, United Kingdom*
[5]*The Edmond and Lily Safra Center for Brain Sciences (ELSC), the Hebrew University of Jerusalem, Jerusalem 91120, Israel*
[6]*The Cognitive Science Program, the Hebrew University of Jerusalem, Jerusalem 91120, Israel*

Abstract: Seed-based functional connectivity (FC) of resting-state functional MRI data is a widely used methodology, enabling the identification of functional brain networks in health and disease. Based on signal correlations across the brain, FC measures are highly sensitive to noise. A somewhat neglected source of noise is the fMRI signal attenuation found in cortical regions in close vicinity to sinuses and air cavities, mainly in the orbitofrontal, anterior frontal and inferior temporal cortices. BOLD signal recorded at these regions suffers from dropout due to susceptibility artifacts, resulting in an attenuated signal with reduced signal-to-noise ratio in as many as 10% of cortical voxels. Nevertheless, signal attenuation is largely overlooked during FC analysis. Here we first demonstrate that signal attenuation can significantly influence FC measures by introducing false functional correlations and diminishing existing correlations between brain regions. We then propose a method for the detection and removal of the attenuated signal ("intensity-based masking") by fitting a Gaussian-based model to the signal intensity distribution and calculating an intensity threshold tailored per subject. Finally, we apply our method on real-world data, showing that it diminishes false correlations caused by signal dropout, and significantly improves the ability to detect functional networks in single subjects. Furthermore, we show that our method increases inter-subject similarity in FC, enabling reliable distinction of different

Additional Supporting Information may be found in the online version of this article.

The first two authors contributed equally to this work.

Contract grant sponsor: Israeli National Science Foundation, Grant number 316/15; and the Ministry of Science and Technology of Israel; Contract grant number: 3-10789.
AA is supported by the European Research Council Grant 310809, and The James S. McDonnell Foundation Award 220020284. UH was supported by the European Research Council Starting Grant NeuroCoDec 309865.

*Correspondence to: Michael Peer, Neuropsychiatry Lab, Dept. of Medical Neurobiology, Faculty of Medicine, The Hebrew University of Jerusalem, Jerusalem 91120, Israel.
E-mail: michael.peer@mail.huji.ac.il

Received for publication 20 October 2015; Revised 13 February 2016; Accepted 4 March 2016.
DOI: 10.1002/hbm.23182
Published online 00 Month 2016 in Wiley Online Library (wileyonlinelibrary.com).

functional networks. We propose to include the intensity-based masking method as a common practice in the pre-processing of seed-based functional connectivity analysis, and provide software tools for the computation of intensity-based masks on fMRI data. *Hum Brain Mapp 00:000–000, 2016.*

Key words: resting-state; pre-processing; noise; artifact; susceptibility; dropout

INTRODUCTION

Resting-state fMRI (RSfMRI) has recently become an extensive research field, offering a method different from the classical event-related paradigms to decode functional neuroanatomy, especially at the level of large-scale networks [Biswal et al., 2010; Fox and Raichle, 2007; Gusnard and Raichle, 2001]. Moreover, it has the potential to serve as a clinical tool for cognitive and mental disorders due to its short protocol and lack of need for explicit subject cooperation [Buckner et al., 2008; Zhang and Raichle, 2010; Peer et al., 2014]. Seed-based analysis, based on the correlation between spontaneous fluctuations of the fMRI signal of different brain voxels or regions (termed "functional connectivity", FC), is one of the most commonly used methods in this field. This method, introduced by Biswal and colleagues already in the mid-90s [Biswal et al., 1995], has repeatedly revealed the existence of correlations between areas that are part of the same functional networks [Fox et al., 2005; Golland et al., 2007; Power et al., 2011; Smith et al., 2009]. Notwithstanding the vast amount of studies investigating the RSfMRI signal, it is not yet used in clinical practice as its reliability varies widely. Due to the reliance on correlation estimates, and the lack of repetitions and averaging, resting-state analysis methods have low signal-to-noise ratio (SNR) and are highly sensitive to noise sources such as motion and magnetic field fluctuations [Jo et al., 2010; Power et al., 2012; Satterthwaite et al., 2012]. Such noise is often spread across different parts of the functional image and therefore might introduce false correlations as well as reduce the strength of real correlations [Satterthwaite et al., 2013]. Therefore, a high SNR is crucial for reliable functional connectivity measurements, as even small amounts of noise may significantly affect results.

One major source of noise in the fMRI BOLD signal is signal dropout due to susceptibility artifacts. These artifacts are mainly caused by air in sinuses and brain cavities which generates a magnetic field inhomogeneity in neighboring cortical regions, mostly inferior temporal and orbitofrontal cortices. Magnetic response to air cavities leads to faster spin decay times and attenuated signal in these regions when acquiring T2* MRI images [but not T1, Deichmann et al., 2002; Ojemann et al., 1997; Fig. 1]. Notably, signal attenuation is more severe in stronger magnetic fields [Krasnow et al., 2003; Poser and Norris, 2009], and thus should be considered as technological advances enable MRI machines with stronger magnetic fields. Signal attenuation has been shown to interfere with identification of BOLD responses to cognitive tasks [Devlin et al., 2000; Shum et al., 2013; Winawer et al., 2010]. While signal attenuation artifacts have been tackled using several methods in the acquisition stage [Deichmann et al., 2002; Glover & Law, 2001; Gorno-Tempini et al., 2002; Weiskopf et al., 2006; Wilson et al., 2002; Ojemann et al., 1997], none of them significantly succeeds in completely abolishing the effect.

The attenuation artifacts under consideration are especially important when performing seed-based functional connectivity analysis. In this technique, connectivity matrices and functional connectivity maps are generated by correlating the time-course of the BOLD signal in single voxels, or in regions of interest by averaging of all voxels in the region [Bullmore and Sporns, 2009]. Importantly, regions of interest are usually defined by constructing a sphere around specific brain coordinates or by pre-defined anatomically parcellated schemes (i.e., brain atlases) [Desikan et al., 2006; Power et al., 2011; Tzourio-Mazoyer et al., 2002; Wang et al., 2009]. Since attenuation artifacts do not affect T1 images, the anatomical definition of a region of interest may include areas of highly attenuated signal (Fig. 1). The effect of this artifact is often neither quantified nor accounted for during FC analysis.

To avoid this contamination of the RSfMRI signal, we provide here a new method of "intensity-based masking". We first describe the calculation of an intensity threshold value for the identification of voxels which contain an attenuated signal. We then explore the extent of such voxels, their temporal SNR and their influence on functional connectivity estimates. Finally, we demonstrate the application of an intensity-based mask on the functional data prior to time-course extraction, and its effects on functional connectivity estimates. The application of the intensity-based masking method is shown to improve the specificity and sensitivity of the FC analysis.

Abbreviations

AAL	Automatic anatomical labeling (atlas)
BOLD	Blood oxygen level dependent
CSF	Cerebro-spinal fluid
FC	Functional connectivity
fMRI	Functional magnetic resonance imaging
HO	Harvard-Oxford (atlas)
RSfMRI	Resting-state functional MRI
tSNR	Temporal signal-to-noise ratio

AQ1

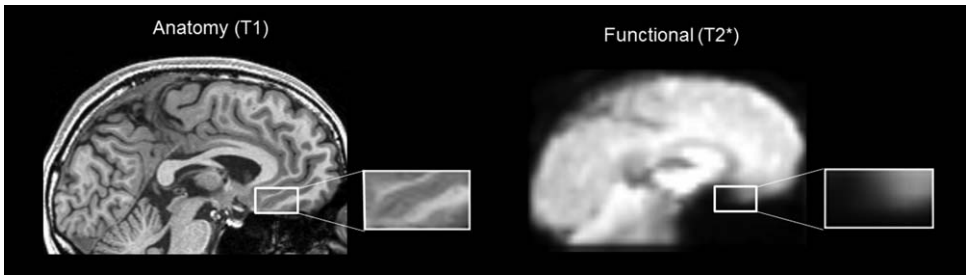

Figure 1.
The signal attenuation (dropout) problem in functional connectivity. Susceptibility artifacts occur only in T2* (functional) MRI images, and not in T1 (anatomical) images. Therefore, a seed region defined on the basis of anatomy may include voxels with strong signal attenuation, which may introduce non-cortical noise to functional time-courses.

MATERIALS AND METHODS

Subjects

Twenty healthy subjects (52.8 ± 17 years old; eight males) participated in the study. Subjects had no personal history of neurologic or psychiatric disorders, and had normal structural MRI. All subjects gave written informed consent, and the study was approved by the ethical committee of the Hadassah Hebrew University Medical Center.

MRI Image Acquisition Procedures

Subjects were scanned using a Siemens Trio 3T system (32 channel head coil). Blood oxygen level dependent (BOLD) fMRI was acquired using a whole-brain, gradient-echo (GE) echo-planar (EPI) sequence of 160 volumes (TR/TE = 2000/30ms, flip angle = 90°, FOV = 192 × 192 mm, matrix = 64 × 64, 33 axial slices, slice thickness/gap = 4 mm/0 mm, voxel size = 3 × 3 × 4 mm). Subjects were instructed to stay awake, keep their eyes open, and remain still. In addition, high-resolution (1 × 1 × 1 mm) T_1-weighted anatomical images (MPRAGE) were acquired for spatial normalization to standard atlas space.

Functional Resting-State MRI Preprocessing

Preprocessing was performed using SPM8 (www.fil.ion.ucl.ac.uk/spm), DPARSFA [Chao-Gan and Yu-Feng, 2010] and Matlab® 2012a (Mathworks, inc.) software. The first five volumes were discarded to ensure magnetization equilibrium. All functional time-series were slice-time corrected, motion corrected to the mean functional image using a tri-linear interpolation with six degrees of freedom, co-registered with the anatomical image, normalized to standard anatomical space (Montreal Neurological Institute EPI template, resampling to 3mm cubic voxels), and spatially smoothed (4mm FWHM, isotropic). Additional preprocessing steps included the removal of linear trends to correct for signal drift and filtering with a 0.01–0.15Hz band-pass filter to reduce non-neuronal contributions to spontaneous BOLD fluctuations. In-line with recent concerns regarding the effect of subjects' motion on functional connectivity characteristics [Van Dijk et al., 2012; Power et al., 2012; Satterthwaite et al., 2012], we performed multiple regression of 24 motion parameters [Friston et al., 1996]: 6 rigid-body head motion parameter values – x, y, z translations and rotations, their value at the prior time point, and the 12 corresponding squared values. In addition, motion "spikes" were also included as regressors (identified by frame-wise displacement of 0.5mm [Power et al., 2012]), such that each regressor equaled zero at all time-points and one at the spike time-point, effectively eliminating the data at the spike without further changes to correlation values [Satterthwaite et al., 2013]. Finally, regressors for global mean, white matter and CSF signals were also included as nuisance sources [Satterthwaite et al., 2013]. Anatomical images were segmented using SPM8's new-segment algorithm, and each voxel's tissue type was identified by max tissue probability: Voxels were labeled as white-matter (WM), grey-matter (GM), cerebrospinal fluid (CSF) or outside the brain (where all other probabilities < 0.2).

Intensity-Based Masking (IBM)

To obtain a model of the intensity distribution, we generated a histogram of maximum BOLD intensity values per voxel for each fMRI acquisition (Fig. 2A). The resulting histogram can be modeled as a sum of two Gaussian distributions and a linear function, where the first Gaussian spans low-intensity voxels (signal attenuation or no brain signal), the second spans high-intensity voxels (white- or grey-matter), and the linear function modeling the transition zone between the two Gaussians (Fig. 2B,C). Curve fitting by optimization (fminsearch) was used to implement the model fit to the BOLD intensity histogram. An intensity-based mask separating high- and low-intensity

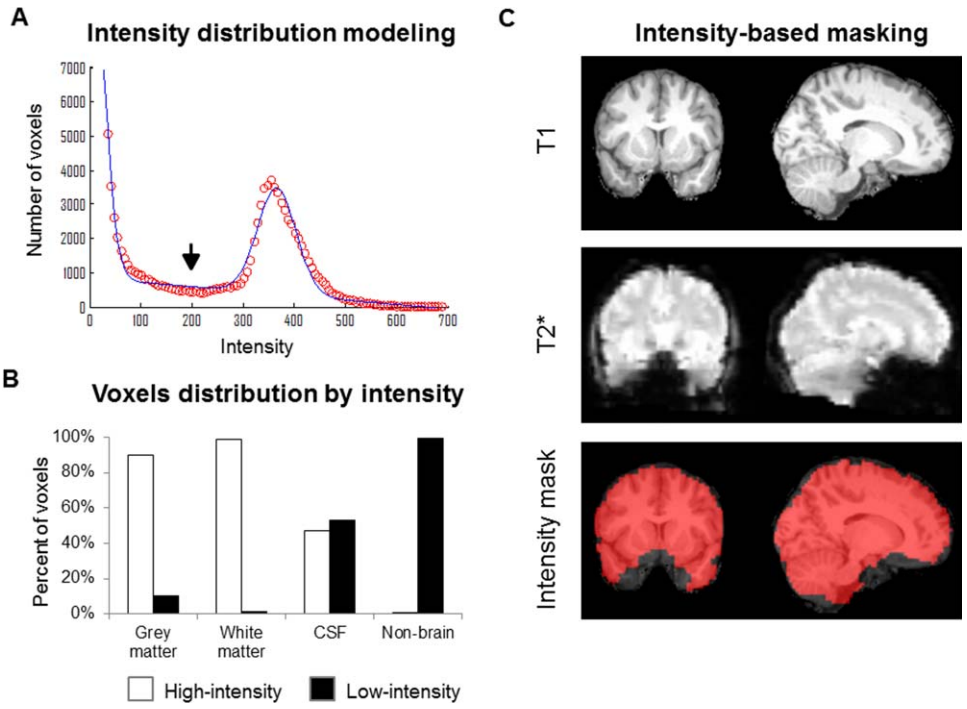

Figure 2.

The *T2** intensity distribution can be modeled as a mixture of Gaussians and a linear trend. (**A**) Intensity distribution (red) and its modeling (blue) as two Gaussians and a linear trend between them, in a representative subject. The intensity threshold (arrow) is defined as the middle point between the peaks of the two Gaussians. (**B**) Percent of high- and low-intensity voxels (above or below the intensity threshold respectively) in grey-matter, white-matter, cerebrospinal fluid (CSF) and outside the brain. White – high-intensity voxels (above intensity threshold), black – low-intensity voxels (below intensity threshold). (**C**) Intensity-based mask created using the computed threshold, overlaid on the brain image of a representative subject. Signal attenuation artifacts are concentrated in the inferior cortical surface (grey regions). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

voxels was created by setting a threshold at the middle point between the peaks of the two Gaussians. Model goodness of fit (R^2) was determined for each subject by squaring the Pearson's correlation coefficient between the optimized model and the intensity distribution. This stage was performed on functional images before application of smoothing, filtering or nuisance-covariates regression (which result in zero-mean voxels). Open-source analysis software for IBM may be downloaded at http://mind.huji.ac.il/intensity_based_masking.aspx.

Extraction of Region-Wise Resting-State Time Series

Brain regions were identified using two commonly used brain atlases: Automatic Anatomical Labeling atlas (AAL), [Tzourio-Mazoyer *et al*. 2002] and Harvard-Oxford atlas (HO) [Desikan et al., 2006]. The atlases were masked with each subject's grey-matter mask. In addition the intensity-based mask derived in the previous stage was applied in order to extract signal from high signal-to-noise voxels only. For each atlas and region, voxels were averaged to produce a time-series of the average activity within that region. Cerebellum regions were not included in the analyses.

Temporal Signal-to-Noise Ratio (tSNR)

tSNR values were calculated for each voxel separately by dividing the voxel's mean signal by its standard deviation across the scan [Murphy et al., 2007]. This stage was performed on the functional images after motion correction, slice-timing correction and normalization, but before smoothing, nuisance covariates regression and filtering (which reduce voxels' mean signal to zero). A two-tailed paired-sample *t*-test was performed to compare between average tSNR values of high-intensity voxels and of low-intensity voxels across subjects. In addition, for each brain atlas tSNR values were averaged across all voxels included in each region (as determined by the specific brain atlas) to compute a regional tSNR value. Correlation was computed between the tSNR value of each atlas region and the number of low-intensity voxels it contains. Signal variance across

time was also computed for each voxel, and its average in low-intensity and high-intensity voxels was compared across subjects using a two-tailed paired-sample *t*-test.

Fourier Transform Analysis

The Fast Fourier Transform (FFT) algorithm (Matlab®, Mathworks inc.) was applied on each atlas region's time-course with and without application of intensity-based masking. Frequencies below 0.1Hz were averaged to calculate the overall amplitude of slow-wave activity in each subject [Biswal et al., 1995]. Amplitudes of slow-wave activity with and without masking were compared using a two-tailed paired-samples *t*-test.

Functional Connectivity

Regional time-courses (obtained either with or without intensity-based masking) were correlated with each other using Pearson's correlation coefficient to produce whole-brain correlation matrices. In addition, regional time-courses were correlated to all the voxels in the brain to produce detailed functional connectivity maps. Normality of the correlation values was obtained using Fisher's *r*-to-z transform. Highly-affected regions were identified as regions with >20% low-intensity voxels, and their connectivity to each other was compared with and without IBM. In addition, connectivity of highly-affected regions to less affected regions (<20% low-intensity voxels) was compared with and without IBM, using only region pairs whose correlation was higher than $r = 0.3$ before IBM application. Additionally, whole-brain correlation matrices were computed for correlation between all possible region pairs using either only low-intensity voxels from each region pair, only high-intensity voxels, or all voxels; these connectivity values were then averaged for each subject across all region pairs, and compared between subjects using two-tailed paired-samples *t*-tests. Furthermore, the difference between the connectivity values with and without masking was computed for each pair of regions per subject; we then computed the correlation between the absolute value of the difference (averaged per region) to the number of the low-intensity voxels the region contains. Visualization was performed using MRIcron [Rorden and Brett, 2000] and Caret [Van Essen et al., 2001].

Connectivity in Literature-Based Resting-State Networks

We used an existing parcellation of the brain into seven resting-state networks (Yeo et al., 2011; available at https://surfer.nmr.mgh.harvard.edu/fswiki/CorticalParcellation_Yeo2011). We next assigned each region in the AAL atlas as related to a specific functional network, by choosing the network with the maximal overlap with it. Correlation values between AAL regions were averaged for all regions belonging to a specific network, for each subject separately, yielding a measure of the average connectivity inside each network. These average network correlations were compared across subjects with and without intensity-based masking using two-tailed paired-samples *t*-tests.

Connectivity Difference Between Subjects

Inter-subject similarity was measured by computing the correlation between their vectorized functional connectivity matrices, for each atlas separately. These similarity measures were computed with and without intensity-based masking, and then compared using two-tailed paired-samples *t*-tests.

Correlation of Regional Signal With Grey-Matter and Non-Brain Signals

The above procedures yielded signal time-courses, extracted from each brain region, and averaged for all low-intensity or high-intensity voxels in each region. We correlated each regional time-course with the average signal from the grey-matter and from the region outside of the brain mask. In addition, time-courses were correlated with each voxel in the image separately, and these correlations were averaged by voxel type as indicated by tissue segmentation (grey-matter or outside-brain voxels). White-matter and CSF voxels' signals were not included in these analyses, as these signals were removed during pre-processing at the nuisance covariates regression stage. Regional average connectivity to grey-matter and non-brain voxels were averaged for each subject (separately for low- and high-intensity voxels), and compared in each subject using two-tailed paired-samples *t*-tests.

RESULTS

Intensity Distribution Modeling and Signal Attenuation Extent

In this section we show the effect of the IBM algorithm on RSfMRI data, and analyze its impact on signal, noise and resulting brain networks. We applied our IBM algorithm fitting a Gaussian-based intensity model (See Materials and methods) to the T2* signal intensity distribution of all subjects (See Fig. 2A for a sample fit; mean R^2-value = 0.98 of fit across subjects). Based on the fitted model, an intensity threshold was calculated for each subject. Gaussian peaks were located at 0 ± 0.1 (noise Gaussian) and 525.5 ± 87.1 (brain Gaussian), and thresholds were therefore set at the middle between the Gaussians at 262.7 ± 43.5 (mean $\pm$ SD across subjects, units are MRI arbitrary intensity units). In the following sections, voxels with intensity values above and below the calculated

threshold (for each individual subject) will be referred to as high- and low-intensity voxels respectively. Unless stated otherwise, the following results will only consider grey-matter voxels.

An examination of the *T*2* images revealed that inside the grey matter, as many as 10.1% of all voxels were low-intensity voxels (below the calculated threshold, averaged across subjects). In particular, an inspection of cortical regions

TABLE I. Extent of signal attenuation artifacts in brain atlases

	Automatic anatomical labeling	Harvard-Oxford
Number of regions with strong signal attenuation artifacts (more than 20% low-intensity voxels)	10/90 (11%)	16/96 (17%)
Maximum percent of low-intensity voxels in an atlas region	57% (gyrus rectus)	56% (subcallosal gyrus)
Atlas regions with strong signal attenuation artifacts (more than 20% low-intensity voxels)	Orbital part of superior frontal gyrus Olfactory cortex Gyrus rectus Parahippocampal gyrus Temporal pole: middle temporal gyrus Inferior temporal gyrus	Temporal pole Inferior temporal gyrus, anterior division Inferior temporal gyrus, posterior division Inferior temporal gyrus, temporooccipital part Frontal medial cortex Subcallosal cortex Frontal orbital cortex Parahippocampal gyrus, anterior division Temporal fusiform cortex, anterior division Temporal fusiform cortex, posterior division

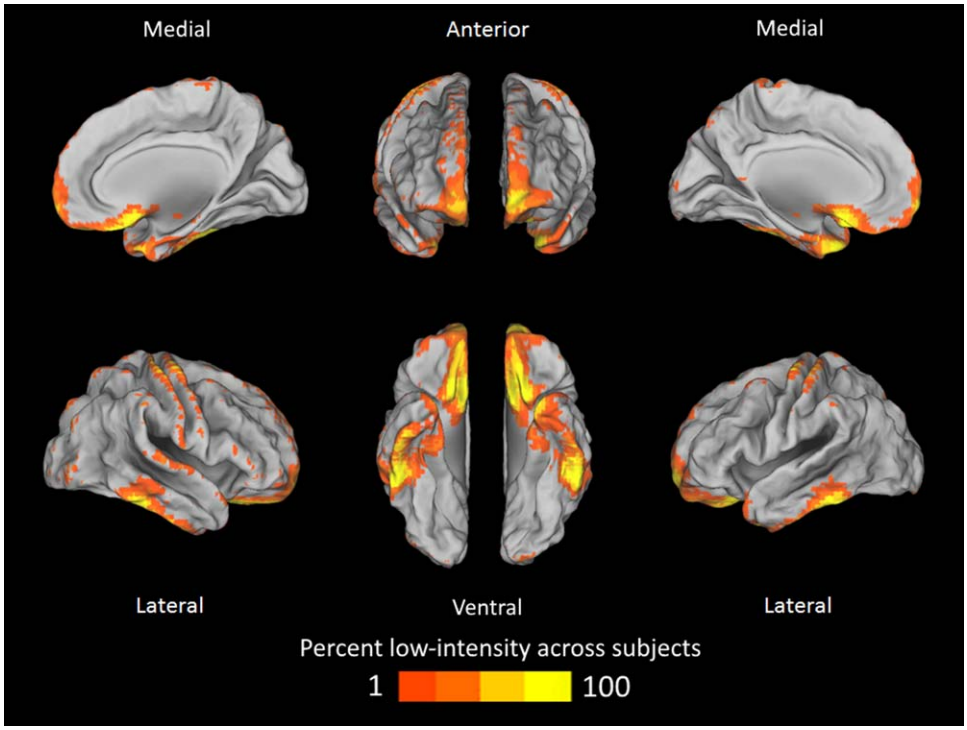

Figure 3.
Whole-brain heat map of low-intensity voxels distribution across subjects. The probability of each voxel to be classified as a low-intensity voxel is shown. Voxels with zero probability are not marked. Low intensity voxels are mainly located in the orbitofrontal, anterior frontal and inferior temporal cortices, and at the edges of gyral ridges, while the medial part of the brain is mostly unaffected.

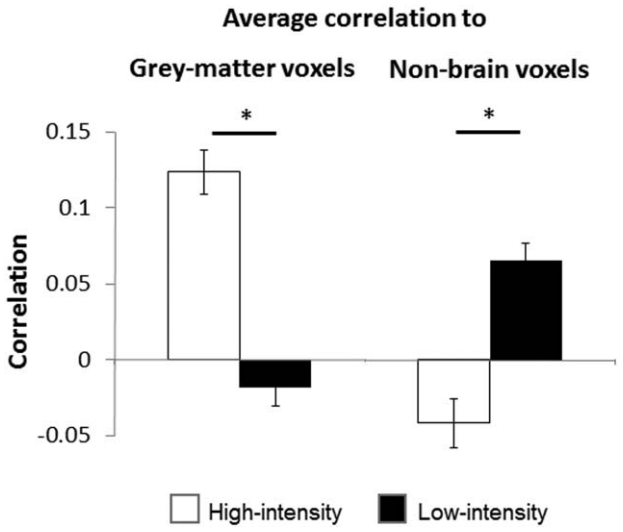

Figure 4.

Correlation of low-intensity and high-intensity grey-matter voxels to other grey-matter voxels and to non-brain regions (voxels outside of the brain). Low-intensity voxels (black) have significantly lower connectivity than high-intensity voxels to grey-matter regions (left) and significantly higher connectivity to non-brain regions (right).

based on the AAL and HO atlases revealed that a large number of atlas regions (11% and 17% respectively, Table I) exhibit particularly strong signal attenuation with >20% of voxels being low-intensity voxels. As expected, the most prominently affected regions were found on the inferior surface of the brain (which is most affected by signal attenuation artifacts), in the orbito-frontal, inferior temporal and anterior frontal cortices, as well as the temporal pole and across the brain at the edges of gyral ridges (Table I, Fig. 3).

Characteristics of Low-Intensity Voxels

Low-intensity voxels were found to have significantly lower temporal signal-to-noise ratio (tSNR) than high-intensity voxels (24.2 and 108.6 respectively, averaged for each subject, paired t-test: $t_{19} = 58.5$, $P < 0.0001$; see Supporting Information Fig. S1 for the tSNR map of a sample subject). In atlas regions, we found a strong and significant negative correlation between each region's tSNR and the number of low-intensity voxels it possesses ($r > 0.78$, $P < 0.0001$ for all atlases). Furthermore, time-courses from low-intensity voxels had a higher signal variance across time than time-courses from high-intensity voxels (67.6 and 40.1 respectively, averaged for each subject, paired t-test: $t_{19} = 8.1$, $P < 0.0001$). This implies that signal attenuation artifacts might account for significant noise in regional signal fluctuations, as low-intensity voxels have a higher per-voxel influence on signal variation across time when averaging data across a region containing both high- and low-intensity voxels.

It is possible, however, that low-intensity voxels may still contain proper, yet attenuated, BOLD signal, and thus their exclusion might result in data loss. To test for the influence of such possibility, we extracted the average time-course of low- and high-intensity voxels from each atlas region. We then checked the correlation, i.e., functional connectivity, between these average time-courses and the average signal from non-brain voxels (voxels outside the brain) and brain voxels (grey-matter voxels). A significantly stronger correlation was found between the signal of low-intensity voxels and the signal from non-brain voxels, compared to the correlation between the signal of high-intensity voxels (in the same regions) and the non-brain voxels (Fig. 4; paired $t_{19} > 4.7$, $P < 0.001$ for both atlases). Furthermore, correlation was significantly weaker between low-intensity voxels and brain-voxels (grey-matter) in comparison with the correlation between high-intensity voxels and brain-voxels (Fig. 4; paired t-test: $t_{19} > 5.5$, $P < 0.0001$ for both atlases). Voxel-wise correlation yielded similar effects both for brain and non-brain voxels ($t_{19} > 2.9$, $P < 0.01$ for all comparisons and atlases). Note that signal from regions outside the brain may be affected by nuisance covariate regression due to shared electronic noise. Nevertheless, the reduced connectivity of low-intensity voxels to the grey matter undermines their contribution to the regional averages of fMRI data, and suggests that their inclusion in seed-based functional connectivity analyses may introduce significant non-brain related noise into the time-course data. Furthermore, as averaging across a region relies on assumptions of homogeneity, voxels with noise will interfere with the overall regional SNR instead of improving it.

We next focused on the connectivity between different atlas-defined regions with large amount (>20%) of low-intensity voxels and all other brain regions, using grey-matter voxels only. When only considering low-intensity voxels, these regions were found to be highly correlated to each other, significantly more than when only considering high-intensity voxels in the same regions pairs (average correlations: AAL: $r = 0.31$, $r = 0.23$, HO: $r = 0.26$, $r = 0.18$, for low and high-intensity voxels respectively, paired t-test: $t_{19} > 3.4$, $P < 0.01$ for both atlases). This suggests that low-intensity voxels are influenced by noise which is correlated across regions. Such contamination can lead to errors in the measurement of FC in prone regions such as the orbitofrontal or lateral temporal cortex. These errors may include both false-positive and false-negative results. For instance, strong false-correlation to close-by regions may eliminate true connection to the important—but more distant—parietal regions (Fig. 5).

Improvement of Overall Functional Connectivity Measures After IBM Application

In order to measure the contribution of our method to FC results, we analyzed our dataset with and without

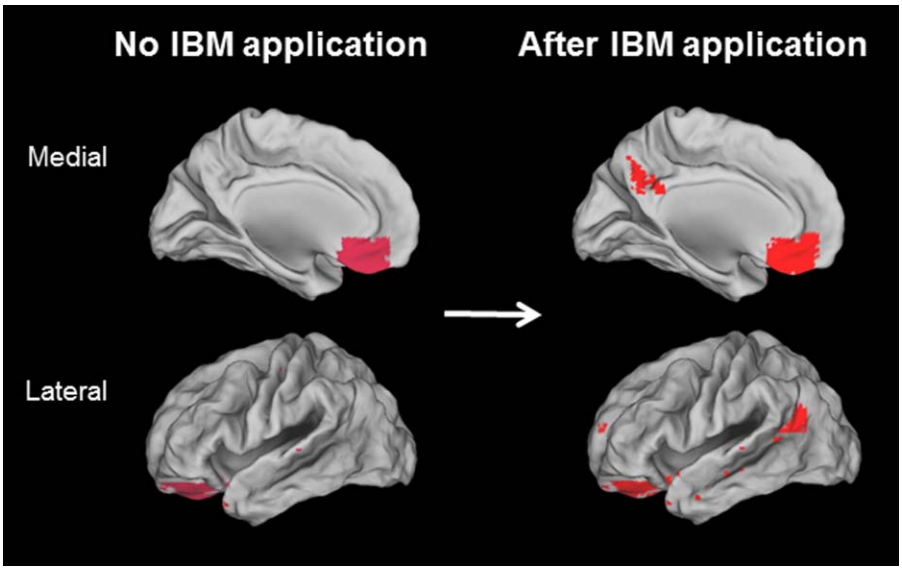

Figure 5.

Improvement of functional connectivity after intensity-based masking. Connectivity of a medial orbitofrontal cortex seed in a representative subject, with and without IBM application. Without IBM, connectivity is confined to nearby regions; following IBM, distant connections are revealed. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

IBM. Removal of low-intensity voxels from each region by masking resulted in significantly lower correlation between regions with a large amount (>20%) of low-intensity voxels (Supporting Information Fig. S2; average correlations: AAL: $r = 0.23$, $r = 0.27$, HO: $r = 0.20$, $r = 0.26$, with and without IBM respectively, paired t-test: $t_{19} > 4.5$, $P < 0.0001$, for both atlases), and on the other hand significantly stronger connections between these regions and regions with small signal loss to which they are connected (average correlations: AAL: $r = 0.48$, $r = 0.43$, HO: $r = 0.44$, $r = 0.41$, with and without IBM respectively, paired t-test: $t_{19} > 6.7$, $P < 0.0001$, for both atlases). Across all regions, a significant correlation was found between the amount of low-intensity voxels removed from a certain region to the absolute difference of its connectivity to the rest of the brain ($r = 0.32$, averaged across atlases and subjects, P -value across subjects < 0.05 for both atlases). Average connectivity between all brain regions had a modest but significant increase across subjects (Average correlation difference: AAL: 0.004, HO: 0.003, paired t-test: $t_{19} > 3.3$, $P < 0.01$ for both atlases), suggesting that IBM removed noise sources which were reducing correlations between brain regions. Notably, in highly affected regions such as the medial orbitofrontal cortex (gyrus rectus) and the subcallosal gyrus, correlations to other brain regions differed by up to 300% after IBM, enabling better detection of functional networks related to these regions (Fig. 5, Supporting Information Fig. S2). In accordance with connectivity increases, Fourier transform analysis indicated that following IBM, regions with low-intensity voxels demonstrated an increase in the amplitude of low frequencies (<0.1 Hz), which has special importance for RSfMRI as it is this frequency which indicates activity of resting-state networks [Biswal et al., 1995] (paired t-test: $t_{19} > 9.3$, $P < 0.0001$ for both atlases).

Next, we investigated how connectivity is affected inside well-known resting-state networks. To this aim we used an existing segregation of the brain into seven functional networks based on connectivity patterns in 1,000 healthy subjects [Yeo et al., 2011]. In five of the seven functional networks (the visual, somatomotor, ventral-attention, fronto-parietal control, and default-mode networks), average connectivity significantly increased after intensity-based masking (paired t-test: $t_{19} > 2.9$, $P < 0.01$ for the five networks). This supports the conclusion that existing intrinsic connectivity patterns are strengthened by application of intensity-based masking. One network – the dorsal-attention network – showed no significant change in connectivity, in accordance with its location in the most dorsal part of the brain (which is unaffected by signal attenuation artifacts). Finally, the inferior orbitofrontal-temporal network showed a significant reduction in connectivity across subjects following intensity-based masking (paired t-test: $t_{19} > 4.5$, $P < 0.0001$), strengthening the conclusion that connectivity between basal brain regions may be largely driven by noise sources in low-intensity voxels.

Finally, we tested similarity in whole-brain connectivity patterns between subjects. In accordance with the above results, correlation between subjects' connectivity matrices increased slightly but significantly (considering all possible

subject pairs, average similarity improvement after IBM: AAL: $r = 0.008$, HO: $r = 0.003$, paired t-test: $t_{189} > 5.4$, $P < 0.0001$ for all atlases), implying that the removal of subject-specific noise sources enables better identification of intrinsic connectivity patterns which are similar across healthy subjects.

DISCUSSION

In this paper we investigated the intensity distribution of the RSfMRI signal and its contamination by attenuation artifacts, and proposed a method for identification of voxels suffering from this artifact by fitting a Gaussian-based model to the intensity distribution. Analyses of low-intensity voxels revealed that these voxels are highly afflicted by noise, resulting in low SNR and high correlation to noise sources outside the brain. We further showed that inclusion of low-intensity voxels can, on the one hand, produce false correlations between functional regions (false-positive), and on the other hand diminish the strength of true connectivity (false-negative). Specifically, our results indicate that connectivity between inferior brain regions may be partially attributed to noise from low-intensity voxels. Finally, application of intensity-based masking (IBM) in order to remove low-intensity voxels prior to FC analysis had led to significant improvement in FC results: better detection of functional networks in each individual subject, lower false connectivity patterns, and increased inter-subject similarity in whole brain connectivity patterns. These effects were especially prominent in connectivity patterns of regions with large signal loss but were also significant at the level of whole-brain connectivity and functional networks.

Our results suggest that IBM should be routinely applied during pre-processing of RSfMRI FC studies, especially when highly-sensitive brain regions (such as the inferior occipital, temporal and orbitofrontal regions) are of interest. Notably, IBM is independent of structural masking (e.g. using a grey-matter mask to obtain only cortical signals), which therefore can and should be simultaneously applied. IBM should be applied regardless of the method used for region-of-interest identification, as all anatomically-based methods (using an anatomically-defined atlas, or defining a seed around specific brain coordinates) cannot account for T2* signal attenuation artifacts. Unlike task-based fMRI where averaging over trials increases SNR, RSfMRI studies should include IBM to eliminate significant contamination. Moreover, signal attenuation artifacts can be subject-specific, suggesting that subject-tailored IBM may be of special importance for understanding of individual subjects or patients. In addition, the IBM method may be helpful for comparisons of groups scanned using different protocols or scanners, as differences may arise in these cases due to differences in the amount of attenuated signal. It may also be beneficial in task-based fMRI if attenuation-sensitive regions are of interest (Abboud et al., 2015). Finally, signal attenuation artifacts are more severe in stronger magnetic fields [Krasnow et al., 2003; Poser and Norris, 2009], and therefore should be taken into account when using MRI scanners with stronger magnetic fields. Despite the importance of IBM, this stage is not in common practice and does not exist in popular brain connectivity analysis tools (such as SPM, DPARSFA, CONN toolbox, AFNI, FSL). Therefore we also provide the needed software tools in order to create intensity-based masks.

The method we propose includes masking away voxels with a high amount of noise. This relies on the fact that to increase SNR, seed-based connectivity analyses average voxels across the defined seed region, based on an assumption of homogeneity of the signal [Yeo et al., 2011; Zang et al., 2004]. Therefore the main question is the contribution of signal vs. noise of each voxel to the average. As we demonstrated, low-intensity voxels contain little signal and large amounts of noise, therefore we recommend their exclusion from the averaging.

Our analysis of the intensity distribution revealed a "transition zone" between high and low intensities (that is, between the Gaussian distributions). An exact threshold inside this transition zone is hard to define. Here we used the distance between the peaks of the two Gaussians to define a threshold which ensures that only high-quality voxels are retained for connectivity calculations; other possible threshold definition methods may be based on the distance from the peak of the first or second Gaussians. Future applications of the IBM methods may help to optimize this parameter.

Previous approaches to dealing with signal attenuation artifacts went in three different directions: improvement of acquisition protocols, increasing post-acquisition SNR by noise removal methods, and correction of correlation coefficients. Improvement of acquisition protocols includes methods such as 3D z-shimming [Glover, 1999], spiral imaging protocols [Glover and Law, 2001; Weiger et al., 2002], and optimization of EPI parameters [Deichmann et al., 2002; Ojemann et al., 1997; Stenger et al., 2000; Weiskopf et al., 2007; Halai et al., 2014]. While these methods increase signal intensity in attenuated areas, they do not entirely eliminate signal attenuation, and may not be suitable for all fMRI experiments. The second approach, post-acquisition removal of noise sources from affected voxels, may be performed using independent components analysis (ICA) [Griffanti et al., 2014; McKeown et al., 1998] or nuisance covariates regression [Lund et al., 2006]. However, ICA usually requires manual identification of noise components and may not remove all noise sources, and nuisance covariate regression may also not completely resolve noise and artefactual correlations. A third approach which was recently suggested is the estimation of each voxel's reliability from its test-retest replicability of correlation patterns to the rest of the brain and correction of correlation coefficients according to the computed

reliability [Mueller et al., 2015]. While this method has many advantages, it differs from IBM in the fact that it requires long scanning times (10-15 minutes) or multiple scans per subject [Mueller et al., 2015], while IBM provides an easily calculable threshold based only on image intensities. Furthermore, voxel reliability across time may be influenced by other factors besides noise, such as difference in connectivity states along the scanning period [Allen et al., 2014]. Finally, since averaging across a whole ROI relies on the assumption of homogeneity of the underlying signal between all ROI voxels, removal of unreliable voxels provides better estimation of the underlying signal than including them. Nevertheless, when investigating voxel-wise connectivity or when investigating fine-grained ROIs from inside low-intensity regions (a scenario where the affected voxels cannot be removed), the approach employed by Mueller et al., may provide an excellent alternative. In general, application of IBM does not exclude co-application of all the other methods described above, and combination of these methods may provide more reliable measurements of functional connectivity in low-signal regions.

Another alternative to discarding voxels with low signal is to increase overall SNR by performing a group analysis. While group analyses may indeed increase SNR and enable better detection of functional networks even in low-SNR voxels, they do not avoid the problem of false correlations, as highly-affected regions are consistently correlated to each other across subjects due to similarity in noise, as was shown above. Moreover, as FC holds promise as a tool for clinical diagnosis of patients, a solution at the single-subject level is crucial. Our method joins a set of other noise removal protocols which may facilitate the development of FC as a clinical diagnostic tool in individual patients [Griffanti et al., 2014; Power et al., 2012; Satterthwaite et al., 2013].

It is important to note that identification of low-intensity voxels (to be removed from connectivity analyses) should be applied on functional data **before any intensity alterations**. Procedures such as signal normalization, filtering or nuisance covariates regression result in alterations of voxel intensity values, so that a threshold cannot be reliably identified. We propose the definition of an intensity-based mask at the pre-processing of RSfMRI data after the stages of slice-timing correction, motion correction and spatial normalization, but before any other alterations of the data (such as filtering or nuisance covariates regression). The identified mask can then be applied on the data after full pre-processing to discard low-intensity voxels from the analysis.

In conclusion, intensity-based masking (IBM) using Gaussian-based modeling may significantly help avoiding both false-positive and false-negative results in functional connectivity studies, improving true connectivity values and distinction between functional networks, and eliminating false connectivity. Application of this method has higher importance in studies which involve brain regions suspected to contain low-intensity signal as well as in studies of individual subjects and patients. As the field of RSfMRI increases dramatically in both number and range of studies, as well as clinical application, the importance of artifact cleaning rises. We therefore suggest that IBM should be incorporated in the routine pre-processing pipeline of RSfMRI studies. Software for IBM calculation may be downloaded at http://mind.huji.ac.il/intensity_based_masking.aspx.

REFERENCES

Abboud S, Maidenbaum S, Dehaene S, Amedi A (2015): A number-form area in the blind. Nat commun 6.

Allen EA, Damaraju E, Plis SM, Erhardt EB, Eichele T, Calhoun VD (2014): Tracking whole-brain connectivity dynamics in the resting state. Cereb Cortex 24:663–676.

Biswal BB, Mennes M, Zuo X-N, Gohel S, Kelly C, Smith SM, Beckmann CF, Adelstein JS, Buckner RL, Colcombe S, Dogonowski A-M, Ernst M, Fair D, Hampson M, Hoptman MJ, Hyde JS, Kiviniemi VJ, Kötter R, Li S-J, Lin C-P, Lowe MJ, Mackay C, Madden DJ, Madsen KH, Margulies DS, Mayberg HS, McMahon K, Monk CS, Mostofsky SH, Nagel BJ, Pekar JJ, Peltier SJ, Petersen SE, Riedl V, Rombouts SARB, Rypma B, Schlaggar BL, Schmidt S, Seidler RD, Siegle GJ, Sorg C, Teng G-J, Veijola J, Villringer A, Walter M, Wang L, Weng X-C, Whitfield-Gabrieli S, Williamson P, Windischberger C, Zang Y-F, Zhang H-Y, Castellanos FX, Milham MP (2010): Toward discovery science of human brain function. Proc Natl Acad Sci U S A 107:4734–4739.

Biswal B, Yetkin FZ, Haughton VM, Hyde JS (1995): Functional connectivity in the motor cortex of resting human brain using echo-planar mri. Magn Reson Med 34:537–541.

Buckner RL, Andrews-Hanna JR, Schacter DL (2008): The brain's default network: Anatomy, function, and relevance to disease. Ann N Y Acad Sci 1124:1–38.

Bullmore E, Sporns O (2009): Complex brain networks: Graph theoretical analysis of structural and functional systems. Nat Rev Neurosci 10:186–198.

Chao-Gan Y, Yu-Feng Z (2010): DPARSF: A MATLAB Toolbox for "Pipeline" Data Analysis of Resting-State fMRI. Front Syst Neurosci 4:13.

Deichmann R, Josephs O, Hutton C, Corfield DR, Turner R (2002): Compensation of susceptibility-induced BOLD sensitivity losses in echo-planar fMRI imaging. Neuroimage 15: 120–135.

Desikan RS, Ségonne F, Fischl B, Quinn BT, Dickerson BC, Blacker D, Buckner RL, Dale AM, Maguire RP, Hyman BT, Albert MS, Killiany RJ (2006): An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. Neuroimage 31:968–980.

Devlin JT, Russell RP, Davis MH, Price CJ, Wilson J, Moss HE, Matthews PM, Tyler LK (2000): Susceptibility-induced loss of signal: Comparing PET and fMRI on a semantic task. Neuroimage 11:589–600.

Fox MD, Raichle ME (2007): Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. Nat Rev Neurosci 8:700–711.

Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, Raichle ME (2005): The human brain is intrinsically organized

into dynamic, anticorrelated functional networks. Proc Natl Acad Sci U S A 102:9673–9678.
Friston KJ, Williams S, Howard R, Frackowiak RSJ, Turner R (1996): Movement-related effects in fMRI time-series. Magn Reson Med 35:346–355.
Glover GH (1999): 3D z-shim method for reduction of susceptibility effects in BOLD fMRI. Magn Reson Med 42:290–299.
Glover GH, Law CS (2001): Spiral-in/out BOLD fMRI for increased SNR and reduced susceptibility artifacts. Magn Reson Med 46:515–522.
Golland Y, Bentin S, Gelbard H, Benjamini Y, Heller R, Nir Y, Hasson U, Malach R (2007): Extrinsic and intrinsic systems in the posterior cortex of the human brain revealed during natural sensory stimulation. Cereb cortex 17:766–777.
Gorno-Tempini ML, Hutton C, Josephs O, Deichmann R, Price C, Turner R (2002): Echo time dependence of BOLD contrast and susceptibility artifacts. Neuroimage 15:136–142.
Griffanti L, Salimi-Khorshidi G, Beckmann CF, Auerbach EJ, Douaud G, Sexton CE, Zsoldos E, Ebmeier KP, Filippini N, Mackay CE, Moeller S, Xu J, Yacoub E, Baselli G, Ugurbil K, Miller KL, Smith SM (2014): ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. Neuroimage 95:232–247.
Gusnard DA, Raichle ME (2001): Searching for a baseline: Functional imaging and the resting human brain. Nat Rev Neurosci 2:685–694.
Halai AD, Welbourne SR, Embleton K, Parkes LM (2014): A comparison of dual gradient-echo and spin-echo fMRI of the inferior temporal lobe. Human brain mapping 35(8):4118–4128.
Jo HJ, Saad ZS, Simmons WK, Milbury LA, Cox RW (2010): Mapping sources of correlation in resting state FMRI, with artifact detection and removal. Neuroimage 52:571–582.
Krasnow B, Tamm L, Greicius MD, Yang TT, Glover GH, Reiss AL, Menon V (2003): Comparison of fMRI activation at 3 and 1.5 T during perceptual, cognitive, and affective processing. Neuroimage 18:813–826.
Lund TE, Madsen KH, Sidaros K, Luo W-L, Nichols TE (2006): Non-white noise in fMRI: Does modelling have an impact? Neuroimage 29:54–66.
McKeown MJ, Makeig S, Brown GG, Jung T-P, Kindermann SS, Bell AJ, Sejnowski TJ (1998): Analysis of fMRI data by blind separation into independent spatial components. Hum Brain Mapp 6:160–188.
Mueller S, Wang D, Fox MD, Pan R, Lu J, Li K, Sun W, Buckner RL, Liu H (2015): Reliability correction for functional connectivity: Theory and implementation. Hum Brain Mapp 36:4664–4680.
Murphy K, Bodurka J, Bandettini PA (2007): How long to scan? The relationship between fMRI temporal signal to noise ratio and necessary scan duration. Neuroimage 34: 565–574.
Ojemann JG, Akbudak E, Snyder AZ, McKinstry RC, Raichle ME, Conturo TE (1997): Anatomic localization and quantitative analysis of gradient refocused echo-planar fMRI susceptibility artifacts. Neuroimage 6:156–167.
Peer M, Nitzan M, Goldberg, I, Katz, J, Gomori, JM, Ben-Hur, T, & Arzy, S (2014): Reversible functional connectivity disturbances during transient global amnesia. Annals of neurology 75(5):634–643.
Poser BA, Norris DG (2009): Investigating the benefits of multi-echo EPI for fMRI at 7 T. Neuroimage 45:1162–1172.
Power JD, Barnes KA, Snyder AZ, Schlaggar BL, Petersen SE (2012): Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. Neuroimage 59:2142–2154.
Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL (2011): Functional network organization of the human brain. Neuron 72:665–678.
Rorden C, Brett M (2000): Stereotaxic display of brain lesions. Behav Neurol 12:191–200.
Satterthwaite TD, Wolf DH, Loughead J, Ruparel K, Elliott MA, Hakonarson H, Gur RC, Gur RE (2012): Impact of in-scanner head motion on multiple measures of functional connectivity: Relevance for studies of neurodevelopment in youth. Neuroimage 60:623–632.
Satterthwaite TD, Elliott MA, Gerraty RT, Ruparel K, Loughead J, Calkins ME, Eickhoff SB, Hakonarson H, Gur RC, Gur RE, Wolf DH (2013): An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. Neuroimage 64:240–256.
Shum J, Hermes D, Foster BL, Dastjerdi M, Rangarajan V, Winawer J, Miller KJ, Parvizi J (2013): A brain area for visual numerals. J Neurosci 33:6709–6715.
Smith SM, Fox PT, Miller KL, Glahn DC, Fox PM, Mackay CE, Filippini N, Watkins KE, Toro R, Laird AR (2009): Correspondence of the brain's functional architecture during activation and rest. Proc Natl Acad Sci U S A 106:13040–13045.
Stenger VA, Boada FE, Noll DC (2000): Three-dimensional tailored RF pulses for the reduction of susceptibility artifacts in T2*-weighted functional MRI. Magn Reson Med 44:525–531.
Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, Mazoyer B, Joliot M (2002): Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. Neuroimage 15:273–289.
Van Dijk KRA, Sabuncu MR, Buckner RL (2012): The influence of head motion on intrinsic functional connectivity MRI. Neuroimage 59:431–438.
Van Essen DC, Drury HA, Dickson J, Harwell J, Hanlon D, Anderson CH (2001): An integrated software suite for surface-based analyses of cerebral cortex. J Am Med Informatics Assoc 8:443–459.
Wang J, Wang L, Zang Y, Yang H, Tang H, Gong Q, Chen Z, Zhu C, He Y (2009): Parcellation-dependent small-world brain functional networks: A resting-state fMRI study. Hum Brain Mapp 30:1511–1523.
Weiger M, Pruessmann KP, Österbauer R, Börnert P, Boesiger P, Jezzard P (2002): Sensitivity-encoded single-shot spiral imaging for reduced susceptibility artifacts in BOLD fMRI. Magn Reson Med 48:860–866.
Weiskopf N, Hutton C, Josephs O, Deichmann R (2006): Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: A whole-brain analysis at 3 T and 1.5 T. Neuroimage 33:493–504.
Weiskopf N, Hutton C, Josephs O, Turner R, Deichmann R (2007): Optimized EPI for fMRI studies of the orbitofrontal cortex: Compensation of susceptibility-induced gradients in the readout direction. Magn Reson Med 20:39–49.
Wilson JL, Jenkinson M, de Araujo I, Kringelbach ML, Rolls ET, Jezzard P (2002): Fast, fully automated global and local

magnetic field optimization for fMRI of the human brain. Neuroimage 17:967–976.
Winawer J, Horiguchi H, Sayres RA, Amano K, Wandell BA (2010): Mapping hV4 and ventral occipital cortex: The venous eclipse. J Vis 10:1.
Yeo BTT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR, Fischl B, Liu H, Buckner RL (2011): The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J Neurophysiol 106:1125–1165.
Zhang D, Raichle ME (2010): Disease and the brain's dark energy. Nat Rev Neurol 6:15–28.
Zang Y, Jiang T, Lu Y, He Y, Tian L (2004): Regional homogeneity approach to fMRI data analysis. Neuroimage 22:394–400.