Investigating resting-state functional connectivity in the cervical spinal cord at 3T

Falk Eippert[a*], Yazhuo Kong[a*], Anderson M. Winkler[a], Jesper L. Andersson[a], Jürgen Finsterbusch[b], Christian Büchel[b], Jonathan C.W. Brooks[c], and Irene Tracey[a]

[a] Oxford Centre for Functional Magnetic Resonance Imaging of the Brain, Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford, UK

[b] Department of Systems Neuroscience, University Medical Center Hamburg-Eppendorf, Hamburg, Germany

[c] Clinical Research and Imaging Centre, University of Bristol, Bristol, UK

[*] These authors contributed equally.

Corresponding author:

Dr. Falk Eippert

FMRIB Centre

John Radcliffe Hospital

Oxford, OX3 9DU

United Kingdom

Phone: +44 1865 222545 / Email: falk.eippert@ndcn.ox.ac.uk

Abstract

The study of spontaneous fluctuations in the blood-oxygen-level-dependent (BOLD) signal has recently been extended from the brain to the spinal cord. Two ultra-high field functional magnetic resonance imaging (fMRI) studies in humans have provided evidence for reproducible resting-state connectivity between the dorsal horns as well as between the ventral horns, and a study in non-human primates has shown that these resting-state signals are impacted by spinal cord injury. As these studies were carried out at ultra-high field strengths using region-of-interest (ROI) based analyses, we investigated whether such resting-state signals could also be observed at the clinically more prevalent field strength of 3T. In a reanalysis of a sample of 20 healthy human participants who underwent a resting-state fMRI acquisition of the cervical spinal cord, we were able to observe significant dorsal horn connectivity as well as ventral horn connectivity, but no consistent effects for connectivity between dorsal and ventral horns, thus replicating the human 7T results. These effects were not only observable when averaging along the acquired length of the spinal cord, but also when we examined each of the acquired spinal segments separately, suggesting that the technique has good spatial sensitivity at 3T. Finally, we investigated the robustness of these resting-state signals against variations in the analysis pipeline by varying the type of ROI creation, temporal filtering, nuisance regression and connectivity metric. We observed that – apart from the effects of band-pass filtering – ventral horn connectivity showed excellent reproducibility, whereas dorsal horn connectivity showed moderate reproducibility. Together, our results demonstrate that spinal cord resting-state connectivity is a robust and spatially consistent phenomenon that could be a valuable tool for investigating the effects of pathology, disease progression, and treatment response in neurological conditions with a spinal component, such as spinal cord injury.

Introduction

The temporal and spatial organization of intrinsic brain activity is currently a subject of intense research. Functional magnetic resonance imaging (fMRI) studies have shown that spontaneous fluctuations in the blood-oxygen-level-dependent (BOLD) signal are organized into distinct and reproducible resting-state networks, such as the sensorimotor, default-mode, or executive-control networks (Buckner et al., 2013; Fox and Raichle, 2007; Power et al., 2014). With the neurophysiological origin of these resting-state signals becoming more evident (Leopold and Maier, 2012; Schölvinck et al., 2013) and their clinical relevance more appreciated (Fox and Greicius, 2010; Zhang and Raichle, 2010), they are increasingly used to probe the integrity and properties of neural circuits in health and disease.

These organized resting-state fluctuations are not an exclusively cortical phenomenon, but have also been observed in subcortical regions as low as the pons and medulla (Beissner et al., 2014; Bianciardi et al., 2016), raising the question whether they constitute a functional signature of the entire central nervous system and might thus be detectable in the spinal cord as well. However, answering this question is a difficult endeavour, because it is challenging to obtain reliable fMRI data from the spinal cord due to a number of issues (Giove et al., 2004; Stroman et al., 2014; Summers et al., 2010), the most prominent of which are: 1) the spinal cord has a very small cross-sectional area (Fradet et al., 2014; Ko et al., 2004), 2) the detrimental influence of physiological noise from cardiac and respiratory sources is much more prominent in the spinal cord than in the brain (Piché et al., 2009; Verma and Cohen-Adad, 2014), and 3) signal loss and image distortion periodically occur along the spinal cord due to the different magnetic susceptibility of vertebrae and connective tissue (Cooke et al., 2004; Finsterbusch et al., 2012).

Despite these obstacles, a few groups have started investigating spinal cord resting-state functional connectivity. Wei and colleagues (2010), for instance explored resting-state signals in the human spinal cord by using independent component analysis (ICA), and reported that the networks detected at the single-subject level were dominated by signal in the frequency range of the respiratory cycle – thus hindering an unequivocal interpretation with regard to a neuronal origin. Building on this initial finding, two other exploratory ICA-based studies used comprehensive denoising strategies and group-level analyses to demonstrate spatially distinct and reproducible spinal cord resting-state signals that are likely to be of neuronal origin (Kong et al., 2014; San Emeterio Nateras et al., 2016). Barry and colleagues used ultra high field imaging at 7T in combination with a hypothesis-driven region-of-interest (ROI) approach to demonstrate robust and reproducible resting-state functional connectivity in the human spinal cord (Barry et al., 2014, 2016). They observed significant time-course correlations between the ventral horns as well as between the dorsal horns at the group level, but not between ventral and dorsal horns. The clinical significance of such resting-state connectivity was recently

demonstrated in a non-human primate model of spinal cord injury, where a spatially-specific influence of lesions on spinal cord functional connectivity was observed (Chen et al., 2015).

These studies hint at the translational potential of using spinal cord resting-state fMRI signals in clinical situations that involve spinal pathology and where a non-invasive metric of disease progression and treatment response would be welcome; examples include multiple sclerosis, spinal cord compression, spinal cord injury, and chronic pain (Wheeler-Kingshott et al., 2014). However, if resting-state connectivity of the human spinal cord is indeed to be used as a potential biomarker for disease progression or treatment effects, some conditions need to be satisfied. First, we need to be able to successfully acquire these signals at the clinically relevant field strength of 3T, because 7T scanners are currently only available in a small minority of research-focussed departments (less than 50 worldwide; Balchandani and Naidich, 2015). Second, we need to demonstrate that the technique is sensitive enough to show robust results when obtaining data from distinct spinal segments, because spinal pathologies can be very localised (for example in spinal cord compression; Nouri et al., 2015). Finally, we need to show that the results we obtain are robust against variations in the data analysis pipeline, thus ensuring their reproducibility (Goodman et al., 2016).

Here, we evaluated to what extent spinal cord resting-state connectivity can satisfy these conditions by reanalysing a previously published data-set that was acquired at a field strength of 3T and used ICA to explore spinal cord resting-state signals (Kong et al., 2014). First, we used this data-set to test whether we could replicate the previously obtained 7T results (significant connectivity between ventral horns and between dorsal horns after averaging over several segments; Barry et al., 2014). Next, we tested whether these results also held for distinct spinal segments – an approach that has only now become possible with the development of a probabilistic atlas of spinal cord segments in a standard space (Cadotte et al., 2015; Cohen-Adad et al., 2014). Third, we assessed whether the obtained results were stable across variations of our data analysis pipeline, by varying the type of 1) ROI creation, 2) temporal filtering, 3) nuisance regression, and 4) connectivity metric. Together, these tests should allow us to determine whether resting-state connectivity in the human spinal cord might be a useful tool in both basic neuroscience and clinical investigations.

Methods

Participants: This study is based on a re-analysis of the data presented in Kong et al. (2014) and thus contains data from the same 20 healthy male participants (age: 26.5 ±3.9 years). The Ethics Committee of the Medical Board in Hamburg, Germany, approved the study and all participants gave written informed consent.

Data acquisition: Magnetic resonance imaging (MRI) data were acquired in an eyes-open state on a 3T system (Magnetom Trio, Siemens, Erlangen, Germany). fMRI data were collected as the last session in a larger spinal fMRI experiment consisting of two sensory and two motor sessions. During all sessions a white crosshair was shown on the screen, which turned red every 15s; participants were asked to stay as still as possible and movements were limited by using a vacuum cushion. Participants were imaged with a 12-channel head coil combined with a 4-channel neck coil (both receive-only), with the cervical spinal cord centred in the neck coil and positioned at isocenter in the magnet. Functional images were acquired using a T2*-weighted gradient-echo echo-planar imaging (EPI) sequence (repetition time 1890ms, echo time 44ms, flip angle: 80°, field of view: 128x128mm², matrix: 128x128, GRAPPA with a PAT-factor of 2). We acquired 16 transversal slices using a slice thickness of 5mm in order to achieve an adequate signal-to-noise ratio despite our high in-plane resolution (1x1mm²). The resulting target volume covered the spinal cord from the 4th cervical vertebra to the 1st thoracic vertebra – based on probabilistic maps of spinal levels, this volume includes segments C6, C7, C8, and T1 (Cadotte et al., 2015). To minimize sensitivity to flow effects, flow rephasing in the slice direction and spatially-selective saturation pulse superior and inferior to the target volume were used and the images obtained with the individual coil channels were combined with a sum-of-squares algorithm. Furthermore, additional saturation pulses were applied posterior and anterior to the target region, i.e. in the phase-encoding direction, in order to avoid pulsatile blood flow artefacts. Signal dropout due to magnetic field inhomogeneity was minimized by using slice-specific z-shimming (Finsterbusch et al., 2012). The adjustments prior to the functional acquisitions (i.e. shimming) were performed on a manually defined volume of about 35x30x70mm³ covering the target region in the spinal cord. Only the neck-coil was used for acquisition of fMRI data and a total of 250 volumes were acquired for each participant (7.5 minute scanning time). To monitor cardiac and respiratory signals during fMRI data acquisition, participants wore a pulse oximeter and respiratory belt, and physiological data were recorded together with the trigger pulses preceding the acquisition of each volume.

We also acquired high-resolution (1x1x1mm³) T_1-weighted anatomical images using a 3D-MPRAGE sequence (sagittal slice orientation, repetition time 2.3s, echo time 3.5ms, flip angle 9°, inversion time 1.1s, field-of-view 192x240x256mm³). The field of view for this acquisition covered an area that spanned at least from the midbrain to

the second thoracic vertebra in every participant; both the neck coil and the head coil were used for this acquisition.

Data processing: Data were processed using tools from FSL (FMRIB Software Library; http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/; Jenkinson et al., 2012). First, each slice was motion corrected for x- and y-translations using FLIRT (FMRIB's Linear Image Registration Tool; http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FLIRT; Jenkinson et al., 2002); translations in the z-direction and rotations were assumed to be minimal, which was confirmed by visual inspection following motion correction. Note that slice-wise motion correction can outperform volume-wise approaches, because spinal cord displacement varies along the rostro-caudal axis of the spinal cord according to the cardiac (Figley and Stroman, 2007) and respiratory cycle (Verma and Cohen-Adad, 2014).

Next, we used FEAT (FMRI Expert Analysis Tool; http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FEAT) to carry out physiological noise regression and high pass filtering (using a cut-off of 100s) – these two steps were performed simultaneously in order to avoid spectral misspecification (Hallquist et al., 2013). The influence of physiological noise of cardiac and respiratory nature is particularly pronounced in the spinal cord and we thus used PNM (Physiological Noise Modelling; http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/PNM; Brooks et al., 2008) in the context of FEAT to remove these noise sources. PNM is based on the RETROICOR approach (Glover et al., 2000) and removes physiological confounds from motion-corrected data using slice-specific regressors based on the calculated phase for each slice relative to the cardiac and respiratory cycles (see also Kong et al., 2012). In the physiological noise model we used here, cardiac, respiratory and interaction effects were modelled using Fourier series, resulting in a total of 32 regressors. Additional nuisance regressors consisted of a) low frequency cerebrospinal fluid (CSF) signal (extracted from voxels whose variance lay in the top 10 percentile within a region including both the spinal cord and CSF space), b) heart rate (value of the smoothed beats per minute (BPM) trace at the acquisition time for each slice), c) motion correction parameters (x and y translation), and d) a regressor that modelled the colour-change of the cross-hair that was presented on the screen. The obtained residuals from each fMRI scan (i.e. physiological noise corrected and high-pass filtered data) were used for further analysis.

Finally, we brought the residuals of the functional data into a common anatomical space. The registration of functional images to the structural volume was initialised using the scanner sqform transformation. Due to EPI distortion in the fMRI data, there remained residual mismatch between the structural and functional data in some slices following the initial transformation. We therefore applied an additional slice-wise registration procedure (x and y translations) on these data (based on hand-drawn spinal cord masks), which minimised the mismatch, and brought fMRI data into good alignment with each participant's structural volume. We then identified

a participant with minimal anterior-posterior and left-right curvature of the spine, which became the experimental template. Subsequently, each participant's structural image was registered to this template using a two-step procedure: we first used FLIRT (FMRIB's Linear Image Registration Tool; http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FLIRT; Jenkinson and Smith, 2001) with default options and angular search range set to 0 degrees and then employed the resulting transformation matrix as a starting point for FNIRT (FMRIB's Non-Linear Image Registration Tool; http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FNIRT; Andersson et al., 2010); important details include: 40 iterations, changing warp resolution from 10mm to 1mm, bias field modelling resolution of 20mm and weighting lambda of 100, weighting mask covering spinal cord and disks/vertebrae. The sqform, XY translation and non-linear warping transformations were then applied to the residuals of the functional data to bring them into a common anatomical space (resampled at 1x1x1mm).

Data analysis – Aim 1: In order to test our first aim – whether we could replicate the results of the recent ROI-based 7T resting-state reports (Barry et al., 2016, 2014) at our field strength of 3T – we needed to create masks for each of the four grey matter horns (along the length of the spinal cord, with our field of view covering segments C6 to T1). These masks were based on a probabilistic grey matter atlas (Taso et al., 2014) that is integrated with the MNI-Poly-AMU template of the spinal cord (Fonov et al., 2014) and is available within SCT (Spinal Cord Toolbox; https://sourceforge.net/projects/spinalcordtoolbox; Cohen-Adad et al., 2014). We obtained these four masks by 1) thresholding the probabilistic grey matter atlas at 50%, 2) splitting the supra-threshold image into 4 different images (one for each horn), 3) making sure that there was at least a one-voxel gap between dorsal and ventral horns, and 4) making sure that the minimal distance between the ventral horns was equal to the minimal distance between the dorsal horns (which resulted in discarding grey matter voxels that belonged to the central grey matter) so that a different distance would not bias the correlation; all steps were done separately for each slice and at the end all slices were merged together. For time-course extraction and statistical analysis, please see below.

Data analysis – Aim 2: In order to test our second aim – whether spinal cord resting-state connectivity could be observed at single segments at our field strength of 3T – we created masks for each horn that did not span the whole extent of the cord, but were instead limited to single spinal cord segments (C6, C7, C8, and T1 in our case). These masks were created by intersecting the previously obtained horn-specific masks with probabilistic masks defining spinal cord segments (Cadotte et al., 2015), which are also integrated with the MNI-Poly-AMU template and available within SCT. We thresholded the probabilistic spinal segment masks at 50% and

minimally edited them manually (removing any overlap between neighbouring segments).

In addition to creating the masks needed to assess Aim 1 (whole-cord correlations) and Aim 2 (single-segment correlations), we also needed to obtain a mapping between our common anatomical space (in which our normalized structural and functional images reside) and the space in which the probabilistic spinal cord atlases reside and where the masks were defined. In order to do so, we first averaged our individual normalized structural images and then applied a non-rigid registration procedure to the resulting average normalized structural image, using the MNI-Poly-AMU template as a target. This was done using procedures implemented in SCT (for details, see De Leener et al., 2014; Fonov et al., 2014) and resulted in deformation fields describing the mapping between the two spaces, allowing us to bring the masks into our common anatomical space.

For both Aim 1 and Aim 2 (where all the following steps were carried out per spinal segment), we used the following procedures to estimate resting-state connectivity. We 1) obtained the average time-course from each of the four horn masks in each participant, 2) calculated Pearson's correlation coefficients between the time-courses for all four horn masks in each participant, 3) averaged the correlation coefficients from left-dorsal-with-left-ventral and right-dorsal-with-right-ventral correlations (to create an index for within-hemicord dorsal-ventral correlations) as well as left-dorsal-with-right-ventral and right-dorsal-with-left-ventral correlations (to create an index for between-hemicord dorsal-ventral correlations) in each participant, and 4) used non-parametric permutation tests for group-level inference. With regard to this last point, we assessed whether the average across subjects of each of the four horn-to-horn correlations (1: dorsal-dorsal, 2: ventral-ventral, 3: within-hemicord dorsal-ventral, 4: between-hemicord dorsal-ventral) was different than zero using permutation testing as implemented in PALM (Permutation Analysis of Linear Models, http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/PALM; Winkler et al., 2014); we used 10000 sign-flips for each test and report two-tailed family-wise-error (FWE) corrected p-values (adjusted for the four tests performed). To give an insight into the inter-individual variability of the four horn-to-horn correlations, we also report the percentage of participants who show each effect. For the whole-cord analysis (Aim 1) we also compared the strength of the different horn-to-horn correlations (1: dorsal-dorsal vs ventral-ventral, 2: dorsal-dorsal vs within-hemicord dorsal-ventral, 3: dorsal-dorsal vs between-hemicord dorsal-ventral, 4: ventral-ventral vs within-hemicord dorsal-ventral, 5: ventral-ventral vs between-hemicord dorsal-ventral, 6: within-hemicord dorsal-ventral vs between-hemicord dorsal-ventral) using the permutation testing analogue of a paired t-test implemented in PALM; we used 10000 permutations and again report two-tailed FWE-corrected p-values (adjusted for the six tests performed). Please note that we report group-averaged correlation coefficients, which tend to exhibit a small conservative bias, i.e. will slightly underestimate the true correlation (compared to averaging Fisher-z transformed correlation coefficients

and then back-transforming the average to Pearson's r, which will slightly overestimate the true correlation; Clayton and Dunlap, 1987; Corey et al., 1998).

Data analysis – Aim 3: In order to test our third aim – which was to assess how robust the observed resting-state connectivity would be against variations in the analysis pipeline, i.e. how reproducible / robust the results would be – we varied the type of 1) ROI creation, 2) temporal filtering, 3) nuisance regression, and 4) connectivity metric.

- *ROI creation*: We not only used the above-described probabilistic masks for each horn (which we will refer to as PROB [for probabilistic]), but also created masks that consisted of just one voxel for each horn (per slice), located at the x-y centre of gravity of each horn (which we will refer to as COG [for centre of gravity]). These masks were created to assess the effects that differences in ROI creation can have on connectivity estimates (Marrelec and Fransson, 2011; Smith et al., 2011) and to ameliorate several issues that could potentially complicate interpreting the data when using the PROB masks: 1) different number of voxels in ventral horns vs dorsal horns (with COG, we just have one voxel per slice per horn), 2) influence of residual CSF fluctuations (with COG, the masks are further away from the subarachnoid space), 3) influence of signal from large veins at the edge of the cord (Cohen-Adad et al., 2010; with COG, the masks are further away from the cord edge), and 4) signal overlap between dorsal horns and ventral horns (due to the point-spread-function of the BOLD response and there only being a one-voxel gap between dorsal horn and ventral horn masks; with COG, the dorsal and ventral horn masks are more strongly separated from each other).
- *Temporal filtering*: Resting-state data have traditionally been band-pass filtered due to the assumption that connectivity is driven by low-frequency fluctuations. However, this has recently been challenged with the discovery that high frequencies also contain meaningful signal (Chen and Glover, 2015; e.g. Niazy et al., 2011). At 7T, Barry and colleagues (Barry et al., 2016) showed that this also holds true for the spinal cord when they noticed that frequencies above 0.08Hz carried meaningful signal, as for example evidenced by higher reproducibility of spinal cord resting-state correlations across sessions. We therefore evaluated the effects of using only a high-pass filter (with a cut-off of 100s, i.e. 0.01Hz; which we will refer to as HIGH) or using a band-pass filter with a pass-band between 0.01 and 0.08Hz (similar to Barry et al., 2014; which we will refer to as BAND). Note that when only using a high-pass filter, we can obtain signals up to the Nyquist frequency of 0.26Hz.
- *Nuisance regression*: We investigated several different slice-wise nuisance regression options in addition to the previously applied slice-wise PNM (see section "Data processing"). First, we investigated the effect of regressing out

the average white matter (WM) signal per slice, which could help to mitigate residual physiological noise effects as well as time-dependent partial volume effects at the grey matter to white matter boundary due to residual motion (Barry et al., 2014). The white matter signal time-course was obtained from the probabilistic white matter mask (thresholded at 10%) – to minimize partial volume effects with grey matter, we subtracted a dilated version of the probabilistic grey matter mask (thresholded at 50%) from this mask. Second, we investigated the effect of regressing out residual cerebrospinal fluid (CSF) signals – in our original PNM, we used one regressor per slice to capture CSF signals, but this might not be sufficient due to CSF flow not being homogenous within the subarachnoid space (CSF flow occurs in different channels with different time-profiles; Schroth and Klose, 1992a; Henry-Feugeas et al., 2000). We therefore carried out a principal component analysis (PCA) on voxel-wise CSF time-courses (which we obtained from the CSF mask that is part of the MNI-Poly-AMU template) and used the first four principal components per slice as regressors (since there are four different CSF channels); together these 4 components explained almost 50% of the variance. Third, we investigated the effect of regressing out non-spinal (NS) signals, i.e. signals that are clearly non-neuronal in origin (e.g. signals in connective tissue or muscles, remaining vascular signals, wide-spread intensity fluctuations due to swallowing, image artefacts, etc.) but might impact on spinal cord BOLD fluctuations. We therefore carried out a PCA on voxel-wise NS time-courses (which we obtained by 1) combining the MNI-Poly-AMU cord mask and CSF mask, 2) dilating the resulting mask and 3) logically inverting the resulting mask) and used the first ten principal components per slice as regressors (each of which explained at least 1% of the variance). This resulted in a total of 8 possible nuisance regression combinations (1: none, 2: WM, 3: WM+CSF, 4: WM+NS, 5: WM+CSF+NS, 6: CSF, 7: CSF+NS, 8: NS).

- *Connectivity metric*: We not only used Pearson's correlation coefficient as described above (which we will refer to as FULL, for full correlation, and which was used by Barry et al., 2014), but also used partial correlation and a regularized version of partial correlation. Partial correlation (which we will refer to as PARTIAL and which was used by Barry et al., 2016) estimates the correlation between two ROIs while controlling for the influence of the time-courses in the remaining two ROIs that do not enter the correlation, i.e. when assessing dorsal-dorsal connectivity this controls for any contributions from the ventral ROIs. This is an attractive approach that is not only able to distinguish between direct and indirect connections (Marrelec et al., 2006), but should also remove any remaining global signal fluctuations that are shared between the ROIs (e.g. residual movement or physiological noise effects). Regularized partial correlation (which we will refer to as REGPARTIAL) imposes a sparseness constraint on the partial correlation matrix and can be beneficial in situations where there are high noise levels, resting-state data

have a short duration, or networks have a large number of ROIs. We used the FSLNETS (http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FSLNets) implementation of regularized partial correlation (which is based on L1-norm regularization, see http://www.cs.ubc.ca/~schmidtm/Software/L1precision.html) with a regularisation-controlling parameter λ of 5 (Smith et al., 2011).

Combining the different factors of ROI creation (two), temporal filtering (two), nuisance regression (eight), and connectivity metric (three) resulted in a total of 96 analyses; for brevity, we only report the results from the whole-cord analyses. In order to gauge the robustness of each of the four horn-to-horn correlations (1: dorsal-dorsal, 2: ventral-ventral, 3: within-hemicord dorsal-ventral, 4: between-hemicord dorsal-ventral) against variations in data analysis, we first investigated whether the sign of the correlation changed with analysis choice, i.e. for each horn-to-horn correlation we report the number of positive correlations among all 96 performed analyses. After observing that only two of the four horn-to-horn correlations were robust against variations in data analysis, we then assessed how the significance of these correlations (again we use FWE-corrected two-tailed p-values as detailed above) was influenced by variations in data analysis, i.e. we report the number of significant correlations among all 96 performed analyses. Supplementing these descriptive reports of the binned data (i.e. positive / negative, significant / not significant), we used a four-way repeated measures analysis of variance (ANOVA; factors ROI creation [2 levels], temporal filtering [2 levels], nuisance regression [8 levels], and connectivity metric [3 levels]) for each of the four horn-to-horn correlation coefficients to investigate the impact of each factor.

Finally, we carried out two complementary analyses that made use of all the 96 different analyses in order to provide evidence for the existence of horn-to-horn connectivity that is identifiable at 3T. In a first analysis we averaged the correlation coefficients across the 96 analyses within each subject and horn-to-horn correlation and then carried out the same permutation test as mentioned above on these averages (again reporting two-tailed FWE-corrected p-values). In a second analysis we used the recently developed modification of non-parametric combination testing (NPC; Winkler et al., 2016) in order to perform joint inference across all 96 analyses using the Fisher combining function. As in the first analysis we report FWE-corrected p-values. Note that these two tests are not equivalent: the null hypothesis for the first test is that the average effect is zero, whereas the null hypothesis for the second test is that all the effects are zero.

Results

Aim 1 – whole cord connectivity: Our first aim was to test whether we could replicate the results of the recent ROI-based 7T resting-state reports (Barry et al., 2016, 2014), namely significant time-course correlations between the two dorsal horns, as well as between the two ventral horns (but not between dorsal and ventral horns) when averaged over the acquired rostro-caudal extent of the spinal cord (Figure 1). Indeed, we observed that the dorsal horns exhibited significant functional connectivity ($r = 0.22$, $p < 0.001$), as did the ventral horns ($r = 0.22$, $p < 0.001$), with 90% of participants showing positive correlations between dorsal horns and 95% of participants showing positive correlations between ventral horns. In contrast to these robust findings, dorsal-ventral horn connectivity within a hemi-cord was just minimally above zero ($r = 0.02$, $p = 0.93$; 50% of participants showed positive correlations), whereas dorsal-ventral horn connectivity between hemicords was significantly negative ($r = -0.08$, $p = 0.03$; 30% of participants showed positive correlations; but see results described under Aim 3).

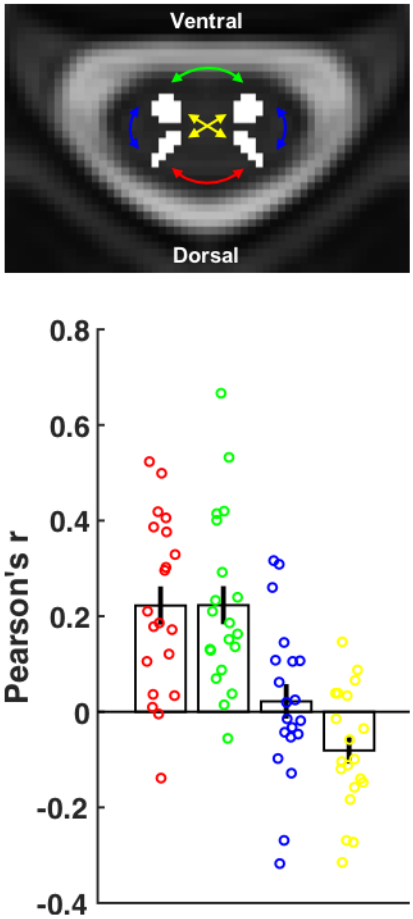

Figure 1. Connectivity averaged along the cord. The transversal slice is taken from the T2-weighted MNI-Poly-AMU template at the middle of segment C6, with the four horn masks overlaid in white and coloured arrows indicating the four different types of horn-to-horn connectivity we investigated (dorsal-dorsal connectivity is depicted in red, ventral-ventral connectivity in green, within-hemicord dorsal-ventral connectivity in blue, and between-hemicord dorsal-ventral connectivity in yellow). The bar-plot displays the group averaged correlation (+/- the standard error of the mean) for each of the four horn-to-horn correlations and the circles indicate participant-specific correlations.

As an aside, we also compared the four different horn-to-horn correlations to each other. We observed that dorsal horn connectivity and ventral horn connectivity were both significantly stronger than a) within-hemicord dorsal-ventral horn connectivity (dorsal: $p = 0.02$, ventral: $p = 0.03$) and b) between-hemicord dorsal-ventral horn connectivity (dorsal: $p < 0.001$, ventral: $p < 0.001$); furthermore, the within-hemicord dorsal-ventral horn connectivity was significantly stronger than the between-hemicord dorsal-ventral horn connectivity ($p = 0.001$; but see results described under Aim 3).

Aim 2 – segment-specific connectivity: Our second aim was to test whether resting-state connectivity could also be observed at single spinal segments (Figure 2). Based on probabilistically defined spinal segments (Cadotte et al., 2015), it is evident that our acquired field of view contains the sixth, seventh and eighth cervical segments as well as the first thoracic segment (C6, C7, C8, T1).

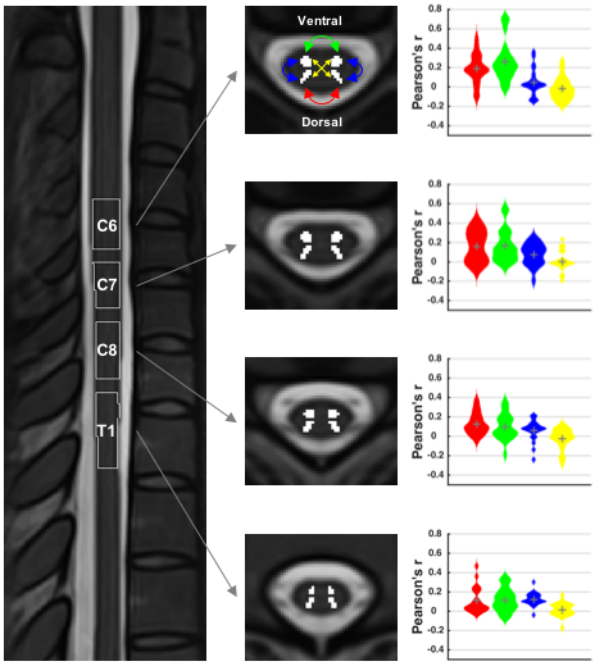

Figure 2: Segment-specific connectivity. The image on the left is a midline sagittal slice through the T2-weighted MNI-Poly-AMU template, with the thresholded probabilistic segments overlaid as outlines. The four transversal slices in the middle are taken from the centre of each of the segments, with the four horn masks overlaid in white. The violin plots on the right demonstrate the correlation between the four horn masks within each segment as smoothed histograms of the distributions (the mean is indicated by the grey plus; color-coding as in Figure 1).

As can be seen in Table 1, the connectivity between dorsal horns as well as between ventral horns was highly significant at every single level, with a minimum of 80% of participants showing positive correlations between dorsal horns and a minimum of 75% of participants showing positive correlations between ventral horns at each spinal level (see also Figure 2). Connectivity between dorsal and ventral horns on the other hand was much more variable, with only some segments (C7 and T1) showing significant results for within-hemicord dorsal-ventral connectivity and none of the segments showing significant results for between-hemicord dorsal-ventral connectivity (Table 1 & Figure 2). These results thus corroborate the whole-cord

connectivity results and show that this technique is sensitive enough to pick up relationships at the level of single spinal segments.

	C6	C7	C8	T1
Dorsal-dorsal				
r	0.19	0.16	0.13	0.11
p	< 0.001	0.001	< 0.001	0.005
%	85	80	95	85
Ventral-ventral				
r	0.26	0.17	0.10	0.12
p	< 0.001	0.001	0.01	0.001
%	95	85	75	80
Within-hemicord				
r	0.05	0.08	0.06	0.12
p	0.35	0.03	0.11	< 0.001
%	70	75	75	95
Between-hemicord				
r	-0.02	0.01	-0.03	0.01
p	0.87	1	0.70	0.81
%	35	40	45	55

Table 1: Segment-specific connectivity. Shown are the results for each of the four spinal levels (C6, C7, C8, and T1) and each of the four correlations (between dorsal horns, between ventral horns, between dorsal and ventral horn within hemicords, between dorsal and ventral horn between hemicords), with *r* representing the average Pearson correlation, *p* representing the two-tailed family-wise-error corrected p-value from a permutation test, and *%* representing the percentage of participants showing a positive correlation.

Aim 3 – robustness of connectivity: Our third aim was to assess how robust the observed resting-state connectivity would be against variations in the analysis pipeline, i.e. how reproducible the results would be. To this end, we carried out a total of 96 analyses using variations of ROI definition (two options), temporal filtering (two options), nuisance regression (eight options), and connectivity metric (three options); for brevity this was only done for the whole-cord connectivity.

First, we observed that the sign of the correlation (i.e. positive or negative connectivity) was remarkably robust against variations in the analysis pipeline for the dorsal horn correlations (96/96 analyses resulted in a positive group-average correlation), as well as for the ventral horn correlations (again, 96/96 analyses resulted in a positive group-average correlation). This was not the case for the correlations between ventral and dorsal horns, where within-hemicord correlations were at chance level (48/96 analyses resulted in a positive group-average correlation) and between-hemicord correlations were also rather variable (64/96 analyses resulted in a positive group-average correlation). When investigating which analysis choices led to this variability, it became clear that for within-hemicord connectivity, ROI creation was the driving factor: positive connectivity was only

observed with the PROB masks and negative connectivity was only observed with the COG masks. This suggests that time-course mixing due to the close proximity of the dorsal and ventral ROIs when using the PROB masks (see Figure 1) is the sole reason for observing positive correlations for within-hemicord dorsal-ventral connectivity, and that these are thus most likely artefactual. For between-hemicord dorsal-ventral connectivity, the influence of analysis choice was not so clear-cut and the only factor we could identify was the use of NS nuisance regression: negative correlations never occurred when this was employed.

Considering that only connectivity between dorsal horns and connectivity between ventral horns seems to be stable across analysis choices, we limited our next analysis – where we assessed whether the significance of the correlations is influenced by variations in the data analysis pipeline – to these two correlations. For the connectivity between dorsal horns, only 24 out of 96 analyses showed a significant correlation, whereas for the connectivity between ventral horns, 65 out of 96 analyses showed a significant correlation (note though that this is a conservative estimate, as FWE-correction was based on four tests). This worryingly low level of statistical robustness could be explained by two factors: temporal filtering and nuisance regression. For the ventral horns, 31 out of the 31 non-significant correlations could be explained by the use of band-pass filtering – within these 31 cases, different variations of nuisance regression were found to contribute as well, though no nuisance regression approach had as much of an impact as band-pass filtering. For the dorsal horns, 43 out of the 72 non-significant correlations could be explained by band-pass filtering – in contrast to the ventral horns nuisance regression had an even stronger impact here with a total of 69 of the 72 non-significant correlations due to variations in this approach, with NS nuisance regression having the strongest impact.

When using a repeated-measures ANOVA to investigate the effects of the different analysis choices on the correlation coefficients more generally (i.e. without binning the data into positive/negative or significant/non-significant), we made three main observations (Table 2): 1) nuisance regression had a consistently strong main effect on connectivity, 2) the main effects of temporal filtering and correlation metric were modest, and 3) ROI creation had a very large main effect on within-hemicord dorsal-ventral connectivity.

Table 2: Effects of variations in data analysis. Shown are the main effects of the four-way repeated-measures ANOVA that was carried out for each of the four horn-to-horn correlations and investigated the impact of different analysis choices.

ROI creation	Temporal filtering	Nuisance regression	Connectivity metric
Dorsal-dorsal			
$F_{1,19} = 3.0$	$F_{1,19} = 0.0$	$F_{7,133} = 13.7$	$F_{1.2,23.3} = 6.8$
p-value = 0.10	p-value = 0.93	p-value < 0.001	p-value = 0.01
Ventral-ventral			
$F_{1,19} = 4.8$	$F_{1,19} = 2.9$	$F_{7,133} = 14.5$	$F_{1.1,20.6} = 2.3$
p-value = 0.04	p-value = 0.10	p-value < 0.001	p-value = 0.15
Within-hemicord			
$F_{1,19} = 184.5$	$F_{1,19} = 1.6$	$F_{7,133} = 1.8$	$F_{1.1,20.4} = 3.0$
p-value < 0.001	p-value = 0.22	p-value = 0.10	p-value = 0.10
Between-hemicord			
$F_{1,19} = 2.3$	$F_{1,19} = 0.0$	$F_{7,133} = 10.3$	$F_{1.0,10.3} = 0.7$
p-value = 0.15	p-value = 0.88	p-value < 0.001	p-value = 0.42

Finally, we used two complementary inferential procedures (both of which make use of all 96 analyses) to test the robustness of the horn-to-horn connections. First, when testing whether the average of all 96 analyses was significant (after FWE-correction for four tests), we observed that the dorsal horn as well as the ventral horn connections were significant (dorsal-dorsal: p = 0.037, ventral-ventral: p = 0.001), whereas the dorsal-ventral connections were not (within-hemicord: p = 0.512, between hemicord: p = 1). Second, we used non-parametric combination (NPC) testing, and observed significant effects for the dorsal horn correlations (p = 0.013), the ventral horn correlations (p < 0.001) and the within-hemicord dorsal-ventral horn correlations (p = 0.017), but not for the between-hemicord dorsal-ventral horn correlations (p = 0.239).

Discussion

In this resting-state fMRI study of the human spinal cord, we observed significant functional connectivity between the dorsal horns, as well as between the ventral horns, but not between the dorsal and ventral horns – neither within hemicords nor between hemicords. These effects were not only evident when considering the whole acquired extent of the cord, but also when considering data from single spinal segments. Finally, we could show that functional connectivity between the ventral horns and between the dorsal horns was mostly robust against variations in the data analysis pipeline, highlighting the reproducibility of these effects.

The motivation for this study arose from results obtained by Barry and colleagues (2014, 2016), who – using ultra high field imaging (7T) of the spinal cord – demonstrated reproducible resting-state functional connectivity between the dorsal horns as well as between the ventral horns. Considering that such connectivity measurements could be a useful tool for probing both disease progression and treatment response in neurological disorders with a strong spinal component, we aimed to test whether we could replicate these findings at the field strength of 3T which is much more prevalent in clinical settings. To this end we reanalysed a previously published resting-state data-set acquired at 3T (Kong et al., 2014) – in our previous publication we had used an exploratory ICA-based approach that is however not suitable for investigating connectivity between a-priori defined ROIs. Here we show that despite many technical differences between the study by Barry and colleagues and our study (e.g. field strength, hardware vendor, fMRI protocol, temporal degrees of freedom, cervical segments covered, estimation of connectivity) we can replicate their main findings at 3T: similar to them, we observed significant resting-state functional connectivity between the ventral horns (which are important for motor function), as well as between the dorsal horns (which are important for sensory function). Both of these findings were robust against inter-individual differences, with at least 90% of participants showing positive connectivity. Also mirroring Barry and colleagues' findings, when averaging across the whole cord we did not observe significant correlations between the dorsal and ventral horns, neither when investigating within-hemicord connectivity nor when investigating between-hemicord connectivity. Consequently, when comparing connectivity strengths, we observed that both dorsal-dorsal and ventral-ventral connectivity was significantly stronger than dorsal-ventral connectivity, both within and between hemicords. In contrast to Barry et al. (2014) we did not observe a significant difference between ventral and dorsal connectivity (see also below). In any case, it is reassuring to see that largely similar results were obtained in these studies despite many technical differences – suggesting that resting-state spinal cord fMRI signals are a robust phenomenon, which bodes well for future multi-centre studies (such as currently underway for spinal cord diffusion tensor imaging: Samson et al., 2016).

Both Barry and colleagues' and our findings were obtained when averaging data along the rostro-caudal axis of the spinal cord – in both instances data were averaged across the whole extent of the field of view. While this approach might help in removing noise and detecting functional connectivity patterns, it ignores the segmental structure of the spinal cord (Baron, 2015) and does not address whether it is possible to detect altered connectivity in localized cord regions as might occur for example with spinal cord compression (Nouri et al., 2015). We therefore tested whether we could detect resting-state connectivity within single spinal cord segments. Investigating this issue at the group level has only become possible very recently with the development of probabilistic maps of spinal segments (Cadotte et al., 2015) and their integration with a standard space template of the spinal cord (Fonov et al., 2014), both openly available with the Spinal Cord Toolbox (https://sourceforge.net/projects/spinalcordtoolbox; Cohen-Adad et al., 2014). With this approach we were able to demonstrate that the overall features of group-level resting-state connectivity – significant dorsal horn correlations and ventral horn correlations – were also detectable at every spinal level we investigated (sixth cervical to first thoracic level), though robustness against inter-individual differences was somewhat lower than for the whole-cord analysis. Segment-wise dorsal-ventral connectivity was never apparent between-hemicords and only partly within-hemicords, though this is most likely artefactual (see below). At least for the segments we imaged, we can thus conclude that dorsal-dorsal and ventral-ventral connectivity seems to be a consistent feature, which our resting-state technique has sufficient sensitivity to detect and might thus be employed in studies investigating neurological disorders with localized spinal pathology.

Before spinal-cord resting-state signals can be used as potential biomarkers (Chen et al., 2015) it is important to assess how robust or reproducible they are (Barry et al., 2016). This has come to the forefront more generally in recent years with concerns regarding the reproducibility of published research in the biomedical literature and beyond (Begley and Ioannidis, 2015; Iqbal et al., 2016). We therefore set out to test how robust / reproducible our obtained results are against reasonable and common variations in the data-analysis pipeline. Note that we are not using the term "reproducibility" as it is used in computational science (where it refers to authors making raw data and analysis code available, so that others can recompute the original results; Peng, 2011), but rather in the form of robustness or inferential reproducibility (Goodman et al., 2016) – i.e. we aim to demonstrate that the basic inferences we draw here (there being strong ventral horn and dorsal horn connectivity) are not conditional on a specific analysis path and should thus be immune to biased data-analyses (Head et al., 2015; which have been discussed under terms such as "data-torturing" Mills, 1993; "researcher degrees of freedom" Simmons et al., 2011; and "p-hacking" Simonsohn et al., 2014); for the sake of brevity we limited these analyses to the whole-cord results.

While there is an enormous flexibility of fMRI data analysis pipelines (Carp, 2012), we chose to focus on a few analysis steps that have received attention in resting-

state studies of the brain, namely ROI creation, temporal filtering, nuisance regression and connectivity metric (e.g. Pruim et al., 2015; Smith et al., 2011). We observed that the sign of the correlation was completely immune to changes in data analysis for dorsal horn connectivity as well as for ventral horn connectivity, but not for dorsal-ventral horn connectivity: within-hemicord dorsal-ventral connectivity only became positive when ROIs were used that were very close to each other, which likely led to time-course mixing (due to the spatial point-spread function of the BOLD response) and inflated positive correlations. We then focussed on the significance of dorsal horn connectivity and ventral horn connectivity and observed that in only about 25% of the different permutations of analysis approaches dorsal horn connectivity was significant; for the ventral horns this was higher with around 70%. When we investigated what was driving this worryingly low robustness, we observed that it was partly due to the use of band-pass filtering: when we ignored the analyses employing a band-pass filter, ventral horn connectivity was significant in 100% of the analyses and dorsal horn connectivity was significant in 40% of the analyses. This is interesting in light of recent findings in the spinal cord (Barry et al., 2016) and the brain (Pruim et al., 2015), where it was demonstrated that band-pass filtering had a negative impact on both the detectability and reproducibility of resting-state connectivity. Consistent with this notion, numerous brain imaging studies have shown that meaningful signal is also contained in higher frequencies above the classical cut-off of 0.08Hz (Boubela et al., 2013; Chen and Glover, 2015; Gohel and Biswal, 2015; Niazy et al., 2011), which might hold true for the spinal cord as well (though thorough characterizations using short-TR data are needed). In addition to these descriptive results, we used two significance tests based on data-aggregation across all 96 analyses and observed significant effects for both dorsal-dorsal and ventral-ventral connectivity in both of these tests, supporting the robustness of these connections. Further support for the idea that dorsal-dorsal and ventral-ventral connectivity is not artificially induced by common noise sources is provided by the fact that full correlation and partial correlation analyses showed similar results.

When investigating the effects of analysis variation more generally using a repeated measures ANOVA based on the correlation coefficients, we observed that nuisance regression had the most consistent influence, while the influence of connectivity metric seemed to be negligible (in line with the above-mentioned findings, temporal filtering had an effect on ventral horn correlations and ROI creation had an enormous effect on within-hemicord dorsal-ventral correlations only). It will be important to tease apart the influence of the various nuisance regression approaches, especially when considering that dorsal horn connectivity seems to be quite prone to these analysis-specific influences. In our opinion the reason for this susceptibility to nuisance regression choice is most likely related to the shape and location of the dorsal horns in comparison to the ventral horns: 1) the dorsal horns are more elongated and thinner, making them more susceptible to partial volume effects with white matter and 2) they border the posterior edge of the cord and could thus be susceptible to partial volume effects with CSF. It remains to be seen how one can

obtain the best balance between preserving "true" signal and removing noise, the latter of which is especially important for resting-state data to avoid false positives due to noise-driven spurious correlations (Cole et al., 2010; Murphy et al., 2013) and even more so in the spinal cord due to its higher level of physiological noise compared to the brain (Cohen-Adad et al., 2010; Piché et al., 2009).

Aside from these technical considerations, an obvious question pertains to the neurobiological underpinnings of the observed signals. This has been discussed in detail previously (e.g. with regard to possible influences of central pattern generators on ventral correlations, input from the peripheral nervous system on dorsal correlations, and supra-spinal influences on both of these; see Barry et al., 2014; Eippert and Tracey, 2014; Kong et al., 2014) and we will only briefly discuss possible underlying factors not mentioned previously. With regard to the dorsal horn connectivity, there is some anatomical evidence for primary afferents that also cross to the contralateral side (Culberson et al., 1979; Light and Perl, 1979) and electrophysiological evidence for an interneuronal network that connects the two dorsal horns (Fitzgerald, 1982, 1983). This has been corroborated more recently with a number of studies identifying populations of dorsal commissural interneurons (Petkó and Antal, 2000; Bannatyne et al., 2006), though these mostly focussed on the lumbar spinal cord. With regard to the ventral horns, there is a wealth of literature on commissural interneurons (for review, see e.g. Jankowska, 2008). While most of these investigations occurred in the upper cervical segments or in the lumbar cord, there is also evidence for commissural systems in the lower cervical segments that we investigated (Alstermark and Kümmel, 1990; Soteropoulos et al., 2013). One should also not discount the effect of respiration on the observed ventral horn resting-state connectivity. Respiration is typically treated as a source of physiological noise in spinal fMRI, e.g. due to breathing-induced B_0 shifts (Verma and Cohen-Adad, 2014) and breathing-induced modulation of CSF flow (Schroth and Klose, 1992b). However, the activity of respiratory motoneurons (Lane, 2011; Monteau and Hilaire, 1991) in the ventral horns might actually underlie some of the observed ventral horn resting-state connectivity. While motoneurons innervating the primary expiratory muscles are unlikely to play a role (the abdominal and internal intercostal muscles are generally not recruited during quiet breathing and are furthermore innervated only from thoracic and lumbar segments), the main inspiratory muscle – the diaphragm – is innervated via the phrenic nerve which originates from segments C3 to C5 in humans (Hollinshead and Keswani, 1956; Routal and Pal, 1999). In this regard it is interesting to note that Barry et al. (2014, 2016) – who acquired data from these segments – observed much stronger correlations between the ventral horns than the dorsal horns, whereas we – who acquired data from below these segments – did not observe such a pronounced difference. Also, resting-state connectivity between ventral horns was almost non-existent in monkeys who were mechanically ventilated during anaesthesia (Chen et al., 2015). Such an interpretation could be tested by investigating whether ventral horn connectivity changes during breathing manipulations. Furthermore, even the ventral horn connectivity we observed could

be driven by respiratory factors, because respiratory interneurons as well as respiratory motoneurons innervating the scalene muscles are present in the lower cervical spinal cord (Lane, 2011; Monteau and Hilaire, 1991; see also Wei et al., 2010). The contribution from these neurons is somewhat unclear however, as the function of respiratory interneurons during normal breathing remains to be elucidated and the scalenes are generally considered only accessory respiratory muscles (but see De Troyer and Estenne, 1984), which also show much lower discharge rates than the diaphragm (Saboisky et al., 2007).

Whatever the neurobiological underpinnings of the observed resting-state signals are, it is important to point out several limitations of the present report. *First*, we need to acknowledge that the observed correlations are rather weak (group-average r of less than 0.3 in most cases), which could stem from the lower temporal signal-to-noise ratio (tSNR) of spinal fMRI data due to the inherent limitations in data acquisition from this structure. In this regard it is also worth mentioning that we did not spatially smooth the data: while smoothing will boost the tSNR, it would also introduce a large amount of time-course mixing between the ROIs and was thus omitted. In the future one might consider using more advanced approaches such as smoothing solely along the cord axis (Cohen-Adad et al., 2014) or non-local spatial filtering (Manjón et al., 2010). It is also possible that our denoising approach was not successful in characterising and removing noise sources properly, but we think this to be unlikely considering the variety of (mostly validated) methods we have employed for noise removal. A final – and neurobiologically more interesting - consideration relates to the possibility that the resting-state signal in each horn of a segment is not only determined by inputs from other horns of this segment, but might be strongly influenced by inter-segmental input (see also Kong et al., 2014) as well as input from supraspinal regions and the peripheral nervous system. This relates to a *second* point, namely that we only investigated intra-segmental, but not inter-segmental connectivity, as this would have far exceeded the scope of this report. *Third*, it is currently unclear why we were not able to obtain evidence for dorsal-ventral connectivity when considering that both within-hemicord and between-hemicord connectivity is essential for some sensorimotor functions such as reflexes – while absence of evidence obviously does not imply evidence of absence, it might just be the case that spinal cord resting-state fluctuations do not cycle through the whole anatomical repertoire of connections and that tonic inhibition of such a system might play a role. And *fourth*, while we investigated the issue of reproducibility or robustness against variations in data analysis, it will also be crucially important for future studies to investigate the inter-session reliability of spinal cord resting-state signals over days, weeks and months, before they might be used in clinical settings (Zuo and Xing, 2014).

Conclusions

In this study we have replicated previously obtained 7T resting-state results (Barry et al., 2014) at the more widely available field strength of 3T. We have furthermore shown that ROI-based dorsal horn connectivity as well as ventral horn connectivity was highly significant not only when averaged across the length of the cord, but also at each of the acquired spinal cord segments, speaking for our technique's sensitivity. Finally, we have investigated the – currently hotly debated – issue of reproducibility and observed that the obtained results are mostly robust against variations in data analysis. In our opinion this suggests that functional connectivity might be a methodologically robust tool for investigating basic spinal cord research questions, such as the correspondence between resting-state and task-based connectivity (Cole et al., 2014) in the spinal cord, the integration between spinal and supra-spinal processes in health (Büchel et al., 2014) and disease (Freund et al., 2016), or how tonic pain protocols (Segerdahl et al., 2015) might lead to a change and possibly spread of dorsal horn spontaneous fluctuations. Even more important, this technique could complement current approaches for assessing pathology, disease progression, and treatment response in neurological disorders with a profound spinal cord component, such as spinal cord injury.

Acknowledgements

AMW is supported by the National Research Council of Brazil (CNPq; 211534/2013-7). CB is supported by the DFG (FOR 1328 and SFB 936) and ERC (2010-AdG_20100407). IT is supported by a Wellcome Trust Strategic Award (102645) and the NIHR Oxford Biomedical Research Centre. JCWB was supported by the Medical Research Council UK (G0700238). JLA is supported by a Wellcome Trust Strategic Award (098369/Z/12/Z).

References

Alstermark, B., Kümmel, H., 1990. Transneuronal transport of wheat germ agglutinin conjugated horseradish peroxidase into last order spinal interneurones projecting to acromio- and spinodeltoideus motoneurones in the cat. 1. Location of labelled interneurones and influence of synaptic activity on the transneuronal transport. Exp. Brain Res. 80, 83–95.

Andersson, J., Jenkinson, M., Smith, S., 2010. Non-linear registration, aka spatial normalisation. FMRIB technical report TR07JA2.

Balchandani, P., Naidich, T.P., 2015. Ultra-High-Field MR Neuroimaging. AJNR Am. J. Neuroradiol. 36, 1204–1215. doi:10.3174/ajnr.A4180

Bannatyne, B.A., Edgley, S.A., Hammar, I., Jankowska, E., Maxwell, D.J., 2006. Differential projections of excitatory and inhibitory dorsal horn interneurons relaying information from group II muscle afferents in the cat spinal cord. J. Neurosci. 26, 2871–2880. doi:10.1523/JNEUROSCI.5172-05.2006

Baron, E., 2015. Spinal cord and spinal nerves: gross anatomy, in: Standring, S. (Ed.), Gray's Anatomy: The Anatomical Basis of Clinical Practice. pp. 762–773.

Barry, R.L., Rogers, B.P., Conrad, B.N., Smith, S.A., Gore, J.C., 2016. Reproducibility of resting state spinal cord networks in healthy volunteers at 7 Tesla. NeuroImage 133, 31–40. doi:10.1016/j.neuroimage.2016.02.058

Barry, R.L., Smith, S.A., Dula, A.N., Gore, J.C., 2014. Resting state functional connectivity in the human spinal cord. eLife 3, e02812.

Begley, C.G., Ioannidis, J.P.A., 2015. Reproducibility in science: improving the standard for basic and preclinical research. Circ. Res. 116, 116–126. doi:10.1161/CIRCRESAHA.114.303819

Beissner, F., Schumann, A., Brunn, F., Eisenträger, D., Bär, K.-J., 2014. Advances in functional magnetic resonance imaging of the human brainstem. NeuroImage 86, 91–98. doi:10.1016/j.neuroimage.2013.07.081

Bianciardi, M., Toschi, N., Eichner, C., Polimeni, J.R., Setsompop, K., Brown, E.N., Hämäläinen, M.S., Rosen, B.R., Wald, L.L., 2016. In vivo functional connectome of human brainstem nuclei of the ascending arousal, autonomic, and motor systems by high spatial resolution 7-Tesla fMRI. Magma 29, 451–462. doi:10.1007/s10334-016-0546-3

Boubela, R.N., Kalcher, K., Huf, W., Kronnerwetter, C., Filzmoser, P., Moser, E., 2013. Beyond Noise: Using Temporal ICA to Extract Meaningful Information from High-Frequency fMRI Signal Fluctuations during Rest. Front. Hum. Neurosci. 7, 168. doi:10.3389/fnhum.2013.00168

Brooks, J.C.W., Beckmann, C.F., Miller, K.L., Wise, R.G., Porro, C.A., Tracey, I., Jenkinson, M., 2008. Physiological noise modelling for spinal functional magnetic resonance imaging studies. NeuroImage 39, 680–692. doi:10.1016/j.neuroimage.2007.09.018

Büchel, C., Geuter, S., Sprenger, C., Eippert, F., 2014. Placebo analgesia: a predictive coding perspective. Neuron 81, 1223–1239. doi:10.1016/j.neuron.2014.02.042

Buckner, R.L., Krienen, F.M., Yeo, B.T.T., 2013. Opportunities and limitations of intrinsic functional connectivity MRI. Nat. Neurosci. 16, 832–837. doi:10.1038/nn.3423

Cadotte, D.W., Cadotte, A., Cohen-Adad, J., Fleet, D., Livne, M., Wilson, J.R., Mikulis, D., Nugaeva, N., Fehlings, M.G., 2015. Characterizing the location of spinal and vertebral levels in the human cervical spinal cord. AJNR Am. J. Neuroradiol. 36, 803–810. doi:10.3174/ajnr.A4192

Carp, J., 2012. On the plurality of (methodological) worlds: estimating the analytic flexibility of FMRI experiments. Front. Neurosci. 6, 149. doi:10.3389/fnins.2012.00149

Chen, J.E., Glover, G.H., 2015. BOLD fractional contribution to resting-state functional connectivity above 0.1 Hz. NeuroImage 107, 207–218. doi:10.1016/j.neuroimage.2014.12.012

Chen, L.M., Mishra, A., Yang, P.-F., Wang, F., Gore, J.C., 2015. Injury alters intrinsic functional connectivity within the primate spinal cord. Proc. Natl. Acad. Sci. U. S. A. 112, 5991–5996. doi:10.1073/pnas.1424106112

Clayton, N., Dunlap, W.P., 1987. Averaging correlation coefficients: Should Fisher's z transformation be used? J. Appl. Psychol. 72, 146–148. doi:10.1037/0021-9010.72.1.146

Cohen-Adad, J., De Leener, B., Benhamou, M., Levy, S., Touati, J., Cadotte, D., Fleet, D., Cadotte, A., Fehlings, M., Pelletier Paquette, J.P., Thong, W., Taso, M., Collins, L., Callot, V., Fonov, V., 2014. Spinal Cord Toolbox: an open-source framework for processing spinal cord MRI data. OHBM.

Cohen-Adad, J., Gauthier, C.J., Brooks, J.C.W., Slessarev, M., Han, J., Fisher, J.A., Rossignol, S., Hoge, R.D., 2010. BOLD signal responses to controlled hypercapnia in human spinal cord. NeuroImage 50, 1074–1084. doi:10.1016/j.neuroimage.2009.12.122

Cole, D.M., Smith, S.M., Beckmann, C.F., 2010. Advances and pitfalls in the analysis and interpretation of resting-state FMRI data. Front. Syst. Neurosci. 4, 8. doi:10.3389/fnsys.2010.00008

Cole, M.W., Bassett, D.S., Power, J.D., Braver, T.S., Petersen, S.E., 2014. Intrinsic and task-evoked network architectures of the human brain. Neuron 83, 238–251. doi:10.1016/j.neuron.2014.05.014

Cooke, F.J., Blamire, A.M., Manners, D.N., Styles, P., Rajagopalan, B., 2004. Quantitative proton magnetic resonance spectroscopy of the cervical spinal cord. Magn. Reson. Med. 51, 1122–1128. doi:10.1002/mrm.20084

Corey, D.M., Dunlap, W.P., Burke, M.J., 1998. Averaging Correlations: Expected Values and Bias in Combined Pearson rs and Fisher's z Transformations. J. Gen. Psychol. 125, 245–261. doi:10.1080/00221309809595548

Culberson, J.L., Haines, D.E., Kimmel, D.L., Brown, P.B., 1979. Contralateral projection of primary afferent fibers to mammalian spinal cord. Exp. Neurol. 64, 83–97.

De Leener, B., Kadoury, S., Cohen-Adad, J., 2014. Robust, accurate and fast automatic segmentation of the spinal cord. NeuroImage. doi:10.1016/j.neuroimage.2014.04.051

De Troyer, A., Estenne, M., 1984. Coordination between rib cage muscles and diaphragm during quiet breathing in humans. J. Appl. Physiol. 57, 899–906.

Eippert, F., Tracey, I., 2014. The spinal cord is never at rest. eLife 3, e03811.

Figley, C.R., Stroman, P.W., 2007. Investigation of human cervical and upper thoracic spinal cord motion: implications for imaging spinal cord structure and function. Magn. Reson. Med. 58, 185–189. doi:10.1002/mrm.21260

Finsterbusch, J., Eippert, F., Büchel, C., 2012. Single, slice-specific z-shim gradient pulses improve T2*-weighted imaging of the spinal cord. NeuroImage 59, 2307–2315. doi:10.1016/j.neuroimage.2011.09.038

Fitzgerald, M., 1983. Influences of contralateral nerve and skin stimulation on neurones in the substantia gelatinosa of the rat spinal cord. Neurosci. Lett. 36, 139–143.

Fitzgerald, M., 1982. The contralateral input to the dorsal horn of the spinal cord in the decerebrate spinal rat. Brain Res. 236, 275–287.

Fonov, V.S., Le Troter, A., Taso, M., De Leener, B., Lévêque, G., Benhamou, M., Sdika, M., Benali, H., Pradat, P.-F., Collins, D.L., Callot, V., Cohen-Adad, J., 2014. Framework for integrated MRI average of the spinal cord white and gray matter: The MNI-Poly-AMU template. NeuroImage 102P2, 817–827. doi:10.1016/j.neuroimage.2014.08.057

Fox, M.D., Greicius, M., 2010. Clinical applications of resting state functional connectivity. Front. Syst. Neurosci. 4, 19. doi:10.3389/fnsys.2010.00019

Fox, M.D., Raichle, M.E., 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. Nat. Rev. Neurosci. 8, 700–711. doi:10.1038/nrn2201

Fradet, L., Arnoux, P.-J., Ranjeva, J.-P., Petit, Y., Callot, V., 2014. Morphometrics of the entire human spinal cord and spinal canal measured from in vivo high-resolution anatomical magnetic resonance imaging. Spine 39, E262-269. doi:10.1097/BRS.0000000000000125

Freund, P., Friston, K., Thompson, A.J., Stephan, K.E., Ashburner, J., Bach, D.R., Nagy, Z., Helms, G., Draganski, B., Mohammadi, S., Schwab, M.E., Curt, A., Weiskopf, N., 2016. Embodied neurology: an integrative framework for neurological disorders. Brain 139, 1855–1861. doi:10.1093/brain/aww076

Giove, F., Garreffa, G., Giulietti, G., Mangia, S., Colonnese, C., Maraviglia, B., 2004. Issues about the fMRI of the human spinal cord. Magn. Reson. Imaging 22, 1505–1516. doi:10.1016/j.mri.2004.10.015

Glover, G.H., Li, T.Q., Ress, D., 2000. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. Magn. Reson. Med. 44, 162–167.

Gohel, S.R., Biswal, B.B., 2015. Functional integration between brain regions at rest occurs in multiple-frequency bands. Brain Connect. 5, 23–34. doi:10.1089/brain.2013.0210

Goodman, S.N., Fanelli, D., Ioannidis, J.P.A., 2016. What does research reproducibility mean? Sci. Transl. Med. 8, 341ps12-341ps12. doi:10.1126/scitranslmed.aaf5027

Hallquist, M.N., Hwang, K., Luna, B., 2013. The nuisance of nuisance regression: spectral misspecification in a common approach to resting-state fMRI preprocessing reintroduces noise and obscures functional connectivity. NeuroImage 82, 208–225. doi:10.1016/j.neuroimage.2013.05.116

Head, M.L., Holman, L., Lanfear, R., Kahn, A.T., Jennions, M.D., 2015. The extent and consequences of p-hacking in science. PLoS Biol. 13, e1002106. doi:10.1371/journal.pbio.1002106

Henry-Feugeas, M.C., Idy-Peretti, I., Baledent, O., Poncelet-Didon, A., Zannoli, G., Bittoun, J., Schouman-Claeys, E., 2000. Origin of subarachnoid cerebrospinal fluid pulsations: a phase-contrast MR analysis. Magn. Reson. Imaging 18, 387–395.

Hollinshead, W.H., Keswani, N.H., 1956. Localization of the phrenic nucleus in the spinal cord of man. Anat. Rec. 125, 683–699.

Iqbal, S.A., Wallach, J.D., Khoury, M.J., Schully, S.D., Ioannidis, J.P.A., 2016. Reproducible Research Practices and Transparency across the Biomedical Literature. PLoS Biol. 14, e1002333. doi:10.1371/journal.pbio.1002333

Jankowska, E., 2008. Spinal interneuronal networks in the cat: elementary components. Brain Res. Rev. 57, 46–55. doi:10.1016/j.brainresrev.2007.06.022

Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. NeuroImage 17, 825–841.

Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M., 2012. FSL. NeuroImage 62, 782–790. doi:10.1016/j.neuroimage.2011.09.015

Jenkinson, M., Smith, S., 2001. A global optimisation method for robust affine registration of brain images. Med. Image Anal. 5, 143–156.

Ko, H.-Y., Park, J.H., Shin, Y.B., Baek, S.Y., 2004. Gross quantitative measurements of spinal cord segments in human. Spinal Cord 42, 35–40. doi:10.1038/sj.sc.3101538

Kong, Y., Eippert, F., Beckmann, C.F., Andersson, J., Finsterbusch, J., Büchel, C., Tracey, I., Brooks, J.C.W., 2014. Intrinsically organized resting state networks in the human spinal cord. Proc. Natl. Acad. Sci. U. S. A. doi:10.1073/pnas.1414293111

Kong, Y., Jenkinson, M., Andersson, J., Tracey, I., Brooks, J.C.W., 2012. Assessment of physiological noise modelling methods for functional imaging of the spinal cord. NeuroImage 60, 1538–1549. doi:10.1016/j.neuroimage.2011.11.077

Lane, M.A., 2011. Spinal respiratory motoneurons and interneurons. Respir. Physiol. Neurobiol. 179, 3–13. doi:10.1016/j.resp.2011.07.004

Leopold, D.A., Maier, A., 2012. Ongoing physiological processes in the cerebral cortex. NeuroImage 62, 2190–2200. doi:10.1016/j.neuroimage.2011.10.059

Light, A.R., Perl, E.R., 1979. Reexamination of the dorsal root projection to the spinal dorsal horn including observations on the differential termination of coarse and fine fibers. J. Comp. Neurol. 186, 117–131. doi:10.1002/cne.901860202

Manjón, J.V., Coupé, P., Martí-Bonmatí, L., Collins, D.L., Robles, M., 2010. Adaptive non-local means denoising of MR images with spatially varying noise levels. J. Magn. Reson. Imaging 31, 192–203. doi:10.1002/jmri.22003

Marrelec, G., Fransson, P., 2011. Assessing the influence of different ROI selection strategies on functional connectivity analyses of fMRI data acquired during steady-state conditions. PloS One 6, e14788. doi:10.1371/journal.pone.0014788

Marrelec, G., Krainik, A., Duffau, H., Pélégrini-Issac, M., Lehéricy, S., Doyon, J., Benali, H., 2006. Partial correlation for functional brain interactivity investigation in functional MRI. NeuroImage 32, 228–237. doi:10.1016/j.neuroimage.2005.12.057

Mills, J.L., 1993. Data torturing. N. Engl. J. Med. 329, 1196–1199. doi:10.1056/NEJM199310143291613

Monteau, R., Hilaire, G., 1991. Spinal respiratory motoneurons. Prog. Neurobiol. 37, 83–144.

Murphy, K., Birn, R.M., Bandettini, P.A., 2013. Resting-state fMRI confounds and cleanup. NeuroImage 80, 349–359. doi:10.1016/j.neuroimage.2013.04.001

Niazy, R.K., Xie, J., Miller, K., Beckmann, C.F., Smith, S.M., 2011. Spectral characteristics of resting state networks. Prog. Brain Res. 193, 259–276. doi:10.1016/B978-0-444-53839-0.00017-X

Nouri, A., Tetreault, L., Singh, A., Karadimas, S.K., Fehlings, M.G., 2015. Degenerative Cervical Myelopathy: Epidemiology, Genetics, and Pathogenesis. Spine 40, E675-693. doi:10.1097/BRS.0000000000000913

Peng, R.D., 2011. Reproducible research in computational science. Science 334, 1226–1227. doi:10.1126/science.1213847

Petkó, M., Antal, M., 2000. Propriospinal afferent and efferent connections of the lateral and medial areas of the dorsal horn (laminae I-IV) in the rat lumbar spinal cord. J. Comp. Neurol. 422, 312–325.

Piché, M., Cohen-Adad, J., Nejad, M.K., Perlbarg, V., Xie, G., Beaudoin, G., Benali, H., Rainville, P., 2009. Characterization of cardiac-related noise in fMRI of the cervical spinal cord. Magn. Reson. Imaging 27, 300–310. doi:10.1016/j.mri.2008.07.019

Power, J.D., Schlaggar, B.L., Petersen, S.E., 2014. Studying brain organization via spontaneous fMRI signal. Neuron 84, 681–696. doi:10.1016/j.neuron.2014.09.007

Pruim, R.H.R., Mennes, M., Buitelaar, J.K., Beckmann, C.F., 2015. Evaluation of ICA-AROMA and alternative strategies for motion artifact removal in resting state fMRI. NeuroImage. doi:10.1016/j.neuroimage.2015.02.063

Routal, R.V., Pal, G.P., 1999. Location of the phrenic nucleus in the human spinal cord. J. Anat. 195 (Pt 4), 617–621.

Saboisky, J.P., Gorman, R.B., De Troyer, A., Gandevia, S.C., Butler, J.E., 2007. Differential activation among five human inspiratory motoneuron pools during tidal breathing. J. Appl. Physiol. 102, 772–780. doi:10.1152/japplphysiol.00683.2006

Samson, R.S., Lévy, S., Schneider, T., Smith, A.K., Smith, S.A., Cohen-Adad, J., Gandini Wheeler-Kingshott, C.A.M., 2016. ZOOM or Non-ZOOM? Assessing Spinal Cord Diffusion Tensor Imaging Protocols for Multi-Centre Studies. PloS One 11, e0155557. doi:10.1371/journal.pone.0155557

San Emeterio Nateras, O., Yu, F., Muir, E.R., Bazan, C., Franklin, C.G., Li, W., Li, J., Lancaster, J.L., Duong, T.Q., 2016. Intrinsic Resting-State Functional Connectivity in the Human Spinal Cord at 3.0 T. Radiology 279, 262–268. doi:10.1148/radiol.2015150768

Schölvinck, M.L., Leopold, D.A., Brookes, M.J., Khader, P.H., 2013. The contribution of electrophysiology to functional connectivity mapping. NeuroImage 80, 297–306. doi:10.1016/j.neuroimage.2013.04.010

Schroth, G., Klose, U., 1992a. Cerebrospinal fluid flow. I. Physiology of cardiac-related pulsation. Neuroradiology 35, 1–9.

Schroth, G., Klose, U., 1992b. Cerebrospinal fluid flow. II. Physiology of respiration-related pulsations. Neuroradiology 35, 10–15.

Segerdahl, A.R., Mezue, M., Okell, T.W., Farrar, J.T., Tracey, I., 2015. The dorsal posterior insula subserves a fundamental role in human pain. Nat. Neurosci. 18, 499–500. doi:10.1038/nn.3969

Simmons, J.P., Nelson, L.D., Simonsohn, U., 2011. False-positive psychology: undisclosed flexibility in data collection and analysis allows presenting anything as significant. Psychol. Sci. 22, 1359–1366. doi:10.1177/0956797611417632

Simonsohn, U., Nelson, L.D., Simmons, J.P., 2014. P-curve: a key to the file-drawer. J. Exp. Psychol. Gen. 143, 534–547. doi:10.1037/a0033242

Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., Woolrich, M.W., 2011. Network modelling methods for FMRI. NeuroImage 54, 875–891. doi:10.1016/j.neuroimage.2010.08.063

Soteropoulos, D.S., Edgley, S.A., Baker, S.N., 2013. Spinal commissural connections to motoneurons controlling the primate hand and wrist. J. Neurosci. 33, 9614–9625. doi:10.1523/JNEUROSCI.0269-13.2013

Stroman, P.W., Wheeler-Kingshott, C., Bacon, M., Schwab, J.M., Bosma, R., Brooks, J., Cadotte, D., Carlstedt, T., Ciccarelli, O., Cohen-Adad, J., Curt, A., Evangelou, N., Fehlings, M.G., Filippi, M., Kelley, B.J., Kollias, S., Mackay, A., Porro, C.A., Smith, S., Strittmatter, S.M., Summers, P., Tracey, I., 2014. The current state-of-the-art of spinal cord imaging: methods. NeuroImage 84, 1070–1081. doi:10.1016/j.neuroimage.2013.04.124

Summers, P.E., Iannetti, G.D., Porro, C.A., 2010. Functional exploration of the human spinal cord during voluntary movement and somatosensory stimulation. Magn. Reson. Imaging 28, 1216–1224. doi:10.1016/j.mri.2010.05.001

Taso, M., Le Troter, A., Sdika, M., Ranjeva, J.-P., Guye, M., Bernard, M., Callot, V., 2014. Construction of an in vivo human spinal cord atlas based on high-resolution MR images at cervical and thoracic levels: preliminary results. Magma 27, 257–267. doi:10.1007/s10334-013-0403-6

Verma, T., Cohen-Adad, J., 2014. Effect of respiration on the B0 field in the human spinal cord at 3T. Magn. Reson. Med. 72, 1629–1636. doi:10.1002/mrm.25075

Wei, P., Li, J., Gao, F., Ye, D., Zhong, Q., Liu, S., 2010. Resting state networks in human cervical spinal cord observed with fMRI. Eur. J. Appl. Physiol. 108, 265–271. doi:10.1007/s00421-009-1205-4

Wheeler-Kingshott, C.A., Stroman, P.W., Schwab, J.M., Bacon, M., Bosma, R., Brooks, J., Cadotte, D.W., Carlstedt, T., Ciccarelli, O., Cohen-Adad, J., Curt, A., Evangelou, N., Fehlings, M.G., Filippi, M., Kelley, B.J., Kollias, S., Mackay, A., Porro, C.A., Smith, S., Strittmatter, S.M., Summers, P., Thompson, A.J., Tracey, I., 2014. The current state-of-the-art of spinal cord imaging: applications. NeuroImage 84, 1082–1093. doi:10.1016/j.neuroimage.2013.07.014

Winkler, A.M., Ridgway, G.R., Webster, M.A., Smith, S.M., Nichols, T.E., 2014. Permutation inference for the general linear model. NeuroImage 92, 381–397. doi:10.1016/j.neuroimage.2014.01.060

Winkler, A.M., Webster, M.A., Brooks, J.C., Tracey, I., Smith, S.M., Nichols, T.E., 2016. Non-parametric combination and related permutation tests for neuroimaging. Hum. Brain Mapp. 37, 1486–1511. doi:10.1002/hbm.23115

Zhang, D., Raichle, M.E., 2010. Disease and the brain's dark energy. Nat. Rev. Neurol. 6, 15–28. doi:10.1038/nrneurol.2009.198

Zuo, X.-N., Xing, X.-X., 2014. Test-retest reliabilities of resting-state FMRI measurements in human brain functional connectomics: a systems neuroscience perspective. Neurosci. Biobehav. Rev. 45, 100–118. doi:10.1016/j.neubiorev.2014.05.009