



Doctoral school in Cognitive and Brain Sciences
XXIV cycle

Ph.D. Dissertation

**Directional relationships between
BOLD activity and autonomic nervous system
fluctuations revealed by fast fMRI acquisition**

PhD candidate: Vittorio Iacovella

Supervisor: Prof. Uri Hasson



Trento, 31/10/2012

1	BOLD AND AUTONOMIC ACTIVITY	11
	SUMMARY	11
1.1	INTRODUCTION	12
1.2	BOLD: A PRODUCT OF A COMPLEX TRANSFER FUNCTION.....	13
1.3	PHYSIOLOGY AS NOISE.....	15
1.4	THE ROLE OF THE AUTONOMIC NERVOUS SYSTEM INDICANTS IN THE CORTICAL ACTIVITY	21
1.5	SUMMARY	28
2	PHYSIOLOGY – RELATED EFFECTS IN THE BOLD SIGNAL AT REST AT 4T...31	
	SUMMARY	31
2.1	INTRODUCTION	32
2.2	METHODS	36
2.2.1	<i>Imaging</i>	36
2.2.2	<i>fMRI pre-processing</i>	37
2.2.3	<i>Construction of the regressors and quantification procedure</i>	38
2.2.4	<i>Permutation method</i>	38
2.3	RESULTS.....	40
2.3.1	<i>Simulations</i>	40
2.3.2	<i>Permutation tests: comparison between cleaning procedures</i>	41
2.3.3	<i>Permutation tests: changes in tSNR</i>	44
2.3.4	<i>Autocorrelation of the regressors</i>	46
2.4	DISCUSSION	50
2.5	CONCLUSIONS.....	53
3	DIRECTIONAL RELATIONSHIPS BETWEEN BOLD ACTIVITY AND AUTONOMIC NERVOUS SYSTEM FLUCTUATIONS REVEALED BY FAST FMRI ACQUISITION	55

SUMMARY	55
3.1 INTRODUCTION	56
3.1.1 <i>Immediate background</i>	56
3.1.2 <i>Design principles</i>	58
3.2 METHODS	61
3.2.1 <i>Considerations underlying determination of imaging protocol</i>	61
3.2.2 <i>fMRI imaging parameters</i>	63
3.2.2.1 Structural scan parameters:	63
3.2.2.2 Fast-TR parameters:	64
3.2.2.3 Slow-TR parameters	64
3.2.3 <i>Participants</i>	64
3.2.4 <i>Procedure</i>	65
3.2.5 <i>Physiological acquisitions</i>	67
3.2.5.1 Extraction of autonomic-related information:.....	68
3.2.5.2 Behavioral effects on the heart speed:.....	70
3.2.5.3 Time-lagged correlations with proxy vectors:.....	70
3.2.5.4 Spectral characterization of smoothed sequences	72
3.2.6 <i>fMRI Data processing</i>	73
3.2.6.1 Slow-TR analysis and ROI extraction:	74
3.2.6.2 tSNR of fast- and slow-TR dataset:.....	76
3.2.6.3 Fast-TR lagged correlation analysis:.....	76
3.2.6.4 Masks and group analysis of the fast-TR protocols:	78
3.2.6.5 Cross-reference with the task related effects.....	79
3.3 RESULTS	80
3.3.1 <i>Behavioral effects on heart speed:</i>	80
3.3.2 <i>Task – induced effects (Slow – TR session)</i>	80

3.3.3	<i>Simulation: apparent synchrony between BOLD and ANS fluctuation due to low sampling rate:</i>	82
3.3.4	<i>tSNR of fast- and slow-TR dataset.</i>	84
3.3.5	<i>Spectral characterization of smoothed sequences</i>	84
3.3.6	<i>BOLD - Autonomic Activity maps:</i>	85
3.3.6.1	<i>Cross – reference with task – related effects</i>	89
3.4	DISCUSSION	90
3.4.1	<i>BOLD - Autonomic Activity regions.</i>	91
3.4.2	<i>Differences between resting-state and steady-state task.</i>	94
3.4.3	<i>Implications for Deactivation theories</i>	96
3.5	CONCLUSIONS	98
4	GENERAL DISCUSSION	99
	SUMMARY	99
4.1	THE PROBLEM: TWO WAYS TO CONSIDER PHYSIOLOGY – RELATED EFFECTS IN THE BOLD SIGNAL	100
4.2	VALIDITY OF RESULTS COMING FROM PHYSIOLOGICAL NOISE CORRECTION	101
4.3	PHYSIOLOGY AS A RELEVANT FACTOR: DIRECTIONAL RELATIONSHIPS BETWEEN BOLD ACTIVITY AND AUTONOMIC NERVOUS SYSTEM FLUCTUATIONS REVEALED BY FAST FMRI ACQUISITION	102
4.4	GENERAL CONTRIBUTIONS	104
5	REFERENCES	109

Abstract

The problem of the relationship between brain function, characterized by functional magnetic resonance imaging, and physiological fluctuations by means of cardiac / respiratory oscillations is one of the most debated topics in the last decade. In recent literature, a great number of studies are found that focus on both practical and conceptual aspects about this topic.

In this work, we start with reviewing two distinct approaches in considering physiology - related sequences with respect to functional magnetic resonance imaging: one treating physiology - related fluctuations as generators of noise, the other considering them as carriers of cognitively relevant information.

In chapter 2 – “Physiology – related effects in the BOLD signal at rest at 4T”, we consider physiological quantities as generators of noise, and discuss conceptual flaws researchers have to face when dealing with data de-noising procedures. We point out that it can be difficult to show that the procedure has achieved its stated aim, i.e. to remove only physiology - related components from the data. As a practical solution, we present a benchmark for assessing whether correction for physiological noise has achieved its stated aim, based on the principle of permutation testing.

In chapter 3 – “Directional relationships between BOLD activity and autonomic nervous system fluctuations revealed by fast fMRI acquisition”, on the other hand, we will consider autonomic indicants derived from physiological time - series as meaningful components of the BOLD signal. There, we describe a FMRI experiment building on this, where the goal

was to localize brain areas whose activity is directionally related to autonomic one, in a top - down modulation fashion.

In chapter 4 we recap the conclusions we found from the two approaches and we summarize the general contributions of our findings. We point out that bringing together the distinct approaches we reviewed lead us to mainly two contributions. On one hand we thought back the validity of almost established procedures in FMRI resting - state pre-processing pipelines. On the other we were able to say something new about general relationship between BOLD and autonomic activity, resting state fluctuations and deactivation theory.

1 BOLD and autonomic activity

Summary

Functional magnetic resonance imaging (fMRI) research has revealed not only important aspects of the neural basis of cognitive and perceptual functions, but also important information on the relation between high-level brain functions and physiology. One of the most relevant questions, given the features of the Blood Oxygenation Level-Dependent (BOLD) signal, is whether and how autonomic nervous system (ANS) functions are related to changes in brain states as measured in the human brain. A straightforward way to address this question has been to acquire external measurements of ANS activity such as cardiac and respiratory data, and examine their relation to the BOLD signal. In this chapter we describe two conceptual approaches to the treatment of ANS measures in the context of BOLD fMRI analysis. On the one hand, several research lines have treated ANS activity measures as noise, considering them as nothing but a confounding factor that reduces the power of fMRI analysis or its validity. Work in this line has developed powerful methods to remove ANS effects from the BOLD signal. On the other hand, a different line of work has made important progress in showing that ANS functions such as cardiac pulsation, heart rate variability and breathing rate, could be considered as a theoretically meaningful component of the signal that is useful for understanding brain function. Work within this latter framework suggests that caution should be exercised when employing procedures to remove correlations between BOLD data and physiological measures. We discuss these two positions and the reasoning underlying them. Following, we draw on the reviewed literature in presenting practical guidelines for treatment of ANS data, which are based on the premise

that ANS data should be considered as theoretically meaningful information. This holds particularly when studying cortical systems involved in regulation, monitoring and/or generation of ANS activity including systems involved in decision making, conflict resolution and the experience of emotion

1.1 Introduction

BOLD functional magnetic resonance imaging (fMRI) is one the most powerful and popular methods for non-invasive examination of whole-brain and regional neural activity patterns. However, the relationship between BOLD fluctuation and pre- or post-synaptic activity at the neural level is non-trivial (Nikos K. Logothetis, 2003). BOLD fluctuations in a given area are induced not only by changes in metabolic demand related to neural activity, but also by a wide range of factors that are considered artifacts such as system noise, hardware limitations (Renvall & Hari, 2009) and confounding effects such as participant's movements inside the scanner. In particular, physiological factors mediating the transfer function between neural and BOLD fluctuations (Nikos K Logothetis, 2008) play an important and potentially confounding factor in the interpretation of BOLD data. Determining the impact of these various noise sources on the BOLD signal is an active domain of research (Bianciardi, Fukunaga, et al., 2009).

One of the main external factors known to co-vary with BOLD measurements is a set of physiological parameters reflecting activity in the autonomic nervous system (ANS). The ANS is a part of the nervous system that controls various functions such as perspiration, respiration, heart rate and blood pressure. It is well documented that fluctuations in these functions are correlated with the BOLD signal and that the impact of such factors may

increase in higher field strengths, rendering measurements performed at higher strength marginally advantageous unless this factor is adequately controlled for (Triantafyllou et al., 2005).

Given the sensitivity of BOLD to ANS functions, one of the important technical and theoretical questions facing scientists relying on BOLD measures is to understand how ANS functions are related to changes in the BOLD signal. To date, this issue has been largely addressed by acquiring external measurements of ANS states and examining their relation to the BOLD signal. With progress in the field, it appears that two implicit approaches to the treatment of ANS data have been developed. On the one hand, studies employing an “ANS as noise” approach have mainly treated physiological acquisitions as artifacts, considering them, from a cognitive or theoretical point of view, as nothing but a confounding factor that reduces the power of an analysis or its validity. Work in this line has developed powerful methods to remove ANS effects from the BOLD signal. On the other hand, other work within various neuroscience domains has made important progress in showing that ANS functions such as cardiac pulsation and breathing rate could be considered as cognitively relevant and interesting components of the BOLD signal. Our aim in this review is to track the rationale for these approaches and derive pragmatic conclusions based on these considerations

1.2 BOLD: a product of a complex transfer function

The BOLD signal is intrinsically a metabolic effect since it is not directly related to electrophysiological changes induced by neuronal populations. Evidence for its biophysical nature was originally shown in PET studies documenting an uncoupling between cerebral blood flow (CBF) and a metabolic quantity of the oxygen consumption ($CMRO_2$) (P. T. Fox, 1986). Quantitative models characterizing aspects of the BOLD effect were developed on the

basis of metabolic variables [e.g. the Balloon Model, (Buxton, Wong, & Frank, 1998)]. The transfer function thought to mediate between neural activity and BOLD has been termed a hemodynamic response function (HRF), to underline both its mechanical and hydrodynamic aspects (Buxton et al., 1998). Numerous following studies explored factors affecting vascular aspects of the BOLD signal, using methods such as the examination of blood flow and volume variations induced by natural or carbon dioxide calibrated respiration. For example, controlled breathing studies have shown the unique impact of vasodilation on the BOLD response (Davis, Kwong, Weisskoff, & Rosen, 1998; Wise, Ide, Poulin, & Tracey, 2004). Other work (Lu, Zhao, Ge, & Lewis-Amezcu, 2008) has shown that factors such as baseline CO₂ levels determine the dynamic range of BOLD and CBF responses to stimulations. One of the interesting points raised by such work is the fact that the asymmetric coupling between CBF and CBV is spatially heterogeneous: this leads to different rates of change in BOLD fluctuations across brain regions. Moreover, inter-individual differences could be well accounted for using such parameters given that individuals' BOLD fluctuations could be related to absolute changes in CO₂ partial pressure, and the final results could be more accurate and specific (Murphy, Harris, & Wise, 2011). In all, these studies have served useful in highlighting how breathing induces changes that affect the BOLD signal. In tandem with studies that have used breathing challenges as a method for understanding the relation between BOLD and physiological factors, other work had demonstrated that although BOLD changes are related to fluctuations in neural activity (N K Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001; Shmuel, Augath, Oeltermann, & Logothetis, 2006), neural activity is not the only factor that causes BOLD fluctuations. In particular, it was recognized that physiology may be one factor affecting ongoing changes in regional metabolism, given that

continuous ANS functions like respiration and cardiac rate trends are tightly linked to the process of oxygenation.

1.3 Physiology as noise

There are a number of ways by which physiological processes can impact BOLD measures, and some of these induce fluctuations that should be treated as noise. For instance, respiration processes may induce systematic motion patterns that reduce the accuracy of BOLD analyses and necessitate correction (Noll & Schneider, n.d.). One type of movement effect is driven by changes in blood pressure through major vessels, which induces movement in nearby tissue (Dagli, Ingeholm, & Haxby, 1999). Respiration induced effects can translate into head motion (Glover, Li, & Ress, 2000) but also result in variations in oxygen concentration and susceptibility effects induced due to the breathing process (Van de Moortele, Pfeuffer, Glover, Ugurbil, & Hu, 2002). Given such impact, work treating ANS correlates as noise has developed analysis methods that treat ANS measurements similarly to how motion data are treated during fMRI preprocessing. A number of approaches have been examined, including the removal of ANS covariates in K-space or during retrospective processing in the image domain (Glover et al., 2000) or the joint removal of physiological effects in tandem with head motion (Jones, Bandettini, & Birn, 2008). Other work has examined related covariates, including heart-rate variance and respiration variance (Chang & Glover, 2009), as well as the effectiveness of various models combining these factors in the context of correction (Jo, Saad, Simmons, Milbury, & Cox, 2010).

In practice, a common general method to link physiological time-courses with fMRI data is the one in which the phase of the physiological time-course is used to produce Fourier expansions of low order. The relation of these expansions with the BOLD data is evaluated

using linear fit against the fMRI time-series. When considering physiological effects as a source of noise, the image-based method for retrospective correction of physiological motion effects (RETROICOR, (Glover et al., 2000)) is commonly utilized to assess the amount of variance in the fMRI data explained by ANS measures. Given that ANS recordings have a higher temporal resolution than BOLD (e.g., 50Hz) some researchers have opted to undersample the measures to the timing of the acquisition TR (van Buuren et al., 2009), whereas others account for ANS effects at the single-slice level, which is acquired at a rate of $[TR / N(\text{slices})]$ Hz.

In absence of physiological recording, several researchers have examined using proxy measures derived from the BOLD signal itself. The particular issue of whether physiological effects are aptly summarized in the “global mean” of the BOLD signal and whether the global mean can be used as a proxy for physiological parameters has been addressed in several studies (M. D. Fox, Zhang, Snyder, & Raichle, 2009; Macey, Macey, Kumar, & Harper, 2004; Schölvinck, Maier, Ye, Duyn, & Leopold, 2010). Furthermore a direct comparison, by means of correlation coefficients, between physiological acquisitions and global mean signal showed that for certain participants it was possible to find a reliable correlation between the global mean of the BOLD response and physiological responses (Chang, Cunningham, & Glover, 2009). However, we note that interpreting the reliability of correlations in this context is complicated by the fact that both the global mean time series and the physiological data themselves demonstrate strong serial autocorrelation, which makes it difficult to assign reliability values to correlations between the two variables, and this is a topic that should be further addressed. The issue of global mean removal has also been examined by several studies, which suggest it is best to not partial out this property for

various analytic reasons (Murphy, Birn, Handwerker, Jones, & Bandettini, 2009), and that at minimum, global mean removal should be employed only in those cases where it does not create interpretive confounds. From a theoretical perspective, recent work (Hyder & Rothman, 2010; Schölvinck et al., 2010), has examined the neuronal correlates of global BOLD fluctuations, by combining local field potentials derived from single-cell recordings and fMRI acquisitions. This work has shown that correlations between gamma frequency (40-80 Hz) of the LFP spectrum and the fMRI signal are observed over large parts of the cortex. This correlation was found to be significant for neural events recorded in diverse brain areas. The fact that about 10% of the fMRI variance could be related to these transient-induced LFP changes suggests that global signal contains cognitively relevant information, and as such should not be removed from the signal. One avenue for future work on removal of physiological effects from the BOLD signal is the development of data driven methods such as ICA to identify components in the BOLD signal whose properties match that of physiological time series (Soldati, Robinson, Persello, Jovicich, & Bruzzone, 2009).

The importance of accounting for potential physiological effects on BOLD data has increased with the theoretical focus on understanding resting state processes in the human and primate brain and their relation to cognition. Since many of the fMRI methods for studying the resting state depend on quantifying patterns of spatially synchronized fluctuations in the BOLD response, any factor that may drive correlated BOLD fluctuations should be corrected for in order to determine that the patterns of spatial correlation identified are indeed driven by synchronized neural sources. This is particularly a concern since physiological measures induce correlation between different brain regions (Chang & Glover, 2009), and given that motion may induce similar patterns of activity that affect the results of

clustering solutions (Mezer, Yovel, Pasternak, Gorfine, & Assaf, 2009) in the resting state (Jones et al., 2008). That is, dissociating between “functional connectivity” and “non-functional connectivity” is becoming an important domain of inquiry. A number of studies have examined the relation between physiological processes, cortical activity, and connectivity amongst various cortical networks.

The relation between respiration effects and the BOLD signal has been evaluated in different ways: modeled as lagged boxcar functions (Kastrup, Li, Glover, & Moseley, 1999), or considered as factors to be convolved with an HRF in a block design context (Thomason, Burrows, Gabrieli, & Glover, 2005). Work by Birn and colleagues examined several interesting aspects of respiration induced effects on the BOLD signal: they started from the observation that variations in respiration rate occur at frequencies (about 0.03 Hz) similar to those related to resting state connectivity (Biswal, Zerrin Yetkin, Haughton, & Hyde, 1995). They showed that the lagged envelope of the respiration time-courses, termed the Respiration Volume per Time (RVT), is an important factor in explaining BOLD fluctuations, both at rest and during a cognitive task (Birn, Diamond, Smith, & Bandettini, 2006). Furthermore, the removal of physiological effects reduced the standard deviation of the signal, suggesting an improvement in data quality. The interrelation between RVT and brain function was examined by probing for correlations between BOLD time-series and RVT lagged regressors obtaining significant values at hemodynamically relevant time lags (about 8 seconds). An extension of this framework lead to the development of a respiration response function (RRF) (Birn, Smith, Jones, & Bandettini, 2008), which was shown to be a valid basis function for respiration induced fluctuations across the brain.

Cardiac effects on BOLD have also been examined. Heart rate and heart rate variability over time have been incorporated into the RETROICOR framework to take in account non-cyclic effects explained by cardio-respiratory phases, and the impact of these factors was evaluated using connectivity studies (van Buuren et al., 2009). Regressors for non-cyclic effects were constructed lagging physiological acquisitions and choosing the shifts that maximized the correlation with BOLD signal. Additional variance (about 5% more than explained by the initial RETROICOR procedure) was found to be explained by non-cyclic effects, including RVT. Importantly, functional connectivity during rest was still reliable after the removal of these effects suggesting that even after physiologically-induced correlations are removed, some correlation is indeed related to neural activity per se.

Chang and colleagues (Chang et al., 2009) aimed to derive a cardiac response function through a deconvolution of the fMRI data. They showed that their model accounted for more variance (taking into account the additional parameter) in about 1/3 of the brain voxels at rest. This improvement was also spatially heterogeneous; it produced an increase of the connectivity within nodes of the default mode network, and a decrease outside it.

The specific issue of the spatial heterogeneity of physiological effects was also examined in several other studies. Saad et al. (Saad et al., 2009) demonstrated that due to the heterogeneity of physiological effects, masks of particular brain areas including white matter and cerebrospinal fluid could be constructed on the basis of structural MPRAGE images, and that the mean time-series from those masks can be used as additional explanatory variables when accounting for physiological noise. A related study demonstrated the use of a completely image-based strategies to remove variance related to global effects (Giove, Gili, Iacovella, Macaluso, & Maraviglia, 2009), and this was shown to improve the data quality of

the signal. This particular attention to tissue features suggests that correction of ANS effects could be better performed when considering the particular anatomical context of various brain regions.

The studies we described explored the relationship between autonomic indicants and neural activity by examining correlations between ANS measures and the BOLD response by examining their relation to connectivity patterns, BOLD data quality or correlation with BOLD response itself. These relations have also been examined outside the domain of BOLD response. For example, optical topography has been used as non-invasive imaging technique to directly evaluate hemoglobin concentrations (HbCC). Katura and colleagues acquired optical topography data as well as two autonomic indicants (HR and mean arterial blood pressure, [MAP]) at rest to examine the relationship between cardiovascular dynamics and low-frequency HbCC fluctuations in the cortex (Katura, Tanaka, Obata, Sato, & Maki, 2006). They evaluated whether there exists a common low-frequency band where oscillations occur both for the hemodynamics and autonomic functions. They identified a frequency band at around 0.1 Hz that could be distinguished within hemodynamic behavior from other fluctuation components. Subsequent work showed that this relation could be modulated pharmacologically (Obrig et al., 2000). Using transfer entropy as a measure of interrelation between HbCC, HR and MAP, Katura et al. also showed that the low-frequency fluctuations in cerebral oxy- and deoxy-hemoglobin could be associated, in a significant fraction (up to 35%), with ANS indicants. This suggests that when analyzing BOLD signal fluctuations in this frequency band, ANS influences could be considered as a significant source of fluctuations.

Work using fMRI has attempted addressing this issue as well. The BOLD signal itself is related to the fluctuations in HbCC, and the blood volume itself could be related to blood pressure effects. Shmueli and colleagues (Shmueli et al., 2007) studied the relationship between cardiac rate variations time-courses and BOLD signal at rest, using correlations and regressions. The ANS indicant they examined was HR variability (HRV), constructed from the cardiac pulsation time-courses. The spectrum of this parameter was shown to be higher in the relevant (about 0.1 Hz) range of frequency. Adding cardiac rate regressors to a RETROICOR framework lead increased variance explained. Furthermore, the spatial distribution of the correlations with the cardiac rate variations was not localized around large vessels (as commonly occurs for cardiac pulsations and respiration time-courses). These results may suggest that cardiac rate variations induce their effects in ways that are not mediated by motion or susceptibility-change effects, and should be considered when one wants to significantly take in account for autonomic nervous system effects on the fluctuations of the BOLD signal.

1.4 The role of the autonomic nervous system indicants in the cortical activity

While the above-reviewed studies aim to characterize the variance explained by various physiological quantities and remove their effect from the BOLD data, another body of work has approached the analysis of ANS indices from a completely different approach, treating these as meaningful signals. This line of work has focused on the neural and behavioral correlates of autonomic indicants in order to understand which brain regions are involved in triggering, modulating or monitoring ANS functions. From the perspective of such studies, the BOLD variance explained by physiological signal is a potentially meaningful component whose properties should be examined.

The idea that bodily states are monitored by the brain and processed at different levels of conscious experience has an established history in the study of psychology and cognitive neuroscience. On an unconscious level, cortical and subcortical structures are known to monitor physiological states such as breathing rate and cardiac responses (Evans, 2010). These monitoring and control operations occur largely on an unconscious level. However, different lines of theoretical and experimental work have also developed the notion that ongoing experience is partly determined by a person's awareness or monitoring of his or her bodily states. One of the earliest and stronger forms of this approach postulated that the emotional aspect of human experience is actually generated as a result of monitoring visceral responses (Lang, 1994). This strong form of the hypothesis was later strongly criticized on various grounds (Cannon, 1987), including the fact that severing connections from internal organs (viscera) to the nervous system still allows the experience of emotions, and that visceral responses may be triggered after emotion is experienced. Recent lesion-based dissociations support this line of criticism, suggesting that the emotion and autonomic responses evoked by emotional stimulus are associated with different systems and are independent of each other (Johnsen, Tranel, Lutgendorf, & Adolphs, 2009).

Nonetheless, the notion that bodily states, including those triggered by the ANS, play a functional role in human experience at different levels has been continuously influential in the cognitive neurosciences. The somatic marker hypothesis developed by Bechara and colleagues (Bechara, 2000) is a theory of human decision making that holds that decisions are "influenced by marker signals that arise in bioregulatory processes, including those that express themselves in emotions and feelings". This approach conjointly holds that bioregulatory processes affect or "colour" ongoing experience as well as decision making

processes. This hypothesis has cited several studies as showing empirical support, most generally for the notion that risky decision making has not only cognitive components but also emotional ones. For instance, whereas non-clinical participants generate skin conductance responses before making an uncertain choice, patients with ventromedial damage do not generate this response. This line of work does not directly address the issue of directionality – i.e., is the relation between ventromedial activity and the ANS response due to an afferent track by which the ventromedial cortex drives ANS activity, or to an efferent track by which the region monitors the ANS state. In addition, given the heterogeneity of this frontal region, it is possible that it is involved in both generation and monitoring of ANS responses.

Subsequent work has documented that instances of conflict resolution, such as those evoked by conflict states (e.g., a Stroop task) similarly induce changes in skin conductance, suggesting that cognitive conflict is associated with ANS changes (Kobayashi, Yoshino, Takahashi, & Nomura, 2007). This has also been documented for cognitive conflict occurring during high level mental processes such as reasoning, which was shown to induce changes in ANS activity (De Neys, Moyens, & Vansteenwegen, 2010). These findings are particularly important given the well established association between one particular ventromedial region, the ACC, and cognitive conflict. To the extent that the ACC is involved in monitoring of somatic states or their generation, such findings could pose an interesting explanation for some ACC activation patterns during conflict.

Consistent with this body of theoretical and experimental work, a body of neuroimaging work to date has attempted to identify neural structures that drive or monitor ANS responses, with particular interest in cortical structures. An important body of work in the area was

pioneered by Critchley and colleagues whose work examines the links between ANS measures and neural function. This work merges classical emotional theories [see for example (Harrison, Gray, Gianaros, & Critchley, 2010)] with physiological and neuroanatomical concepts [see (Craig, 2002, 2009)]. In a recent work (Harrison et al., 2010), cortical correlates of emotional feelings were studied, using fMRI and gastric electrogastrogram (EGG) acquisitions. Participants were instructed to watch and judge videos where several kinds of feelings of disgust were elicited. Changes in ANS indicants were reflected in the activity of the insula: this area showed correlated activity both with increased high frequency heart power. This suggests that the insula could be involved in the representation of afferent information from the ANS towards the cortex (see also (H D Critchley, 2005)). Brain activity increases in the insula and in the inferior parietal lobule were shown to be related to the accuracy in counting heart beats (Pollatos, Schandry, Auer, & Kaufmann, 2007), suggesting the area mediates attention to one's own arousal.

Other support for high-level modulation of the relationship between brain function and ANS was examined in a study by Metz-Lutz and colleagues. In that study fMRI and ECG were acquired to explore human adhesion during movie-watching (the belief that what is shown is real). They studied changes in HRV using a lagged-autocorrelation index (correlation coefficient in Poincarè plots) to explore when heart rate changes occur. They found that decreases in heart-rate variability occurred in the same time as the activation of several brain regions, including Brodmann areas 21/22 and 37/39. From a physiological perspective, this exemplifies how naturally occurring emotions during movie viewing are associated with ANS indicants such as heart rate variability. From a technical viewpoint, such a finding highlights the important of considering this relation during data analysis, since

the link between this ANS measure and cognitive activity implies that procedures that remove the effect of heart rate variability from BOLD signal could well remove meaningful components of the BOLD signal as well.

The relationship between ANS indicants and BOLD has also been examined outside the domain of high-level cognitive function. In an early study (H D Critchley, Corfield, Chandler, Mathias, & Dolan, 2000) PET data were collected while individuals' blood pressure (mean arterial pressure; MAP) and heart rate were monitored. Participants in that study were instructed to perform either effortless or effortful variants of two tasks: an isometric (squeezing a bulb) exercise or a mental arithmetic (serial subtraction) stressor task. Conjunction analyses were performed to explore areas common to both the tasks, revealing activity in subcortical structures including the cerebellar vermis and brainstem, but also in the right anterior cingulate. These same regions and the right insula were shown to covary on a group level with MAP. Finally, activity in the pons, cerebellum, and right insula was associated with heart rate in the same way.

Follow up work used fMRI and electrocardiographic acquisitions (Critchley et al., 2003). In this study, cardiac time-courses were processed to separate low (0.05-0.15 Hz) and high (0.15-0.50 Hz) frequency components that refer, respectively, to sympathetic and parasympathetic neural influences on heart rate (Montano, Porta, & Malliani, 2001). The tasks utilized were similar to those used by Critchley et al. (H D Critchley et al., 2000). The study found a significant correlation between the low frequency cardiac component and anterior cingulate activity. They interpreted this finding as showing an involvement of the anterior cingulate cortex in modulation of cardiac function. Importantly, the same study demonstrated the patients with a lesion in the ACC performed well on both tasks but did not show the same

sort of ANS modulation. Instead, they showed relatively reduced cardiovascular responses to effortful cognitive tasks. This was taken to suggest that the anterior cingulate plays an important role in regulation of bodily states of arousal, to meet concurrent behavioral demands.

The study of the relationship between brain activity and ANS indicants could be performed by modeling their relation to brain activity as mediated by high-level cognitive functions. Wager and colleagues (Wager, Waugh, et al., 2009), demonstrated that social evaluative threat (SET) was correlated with cardiovascular responses. In that study, they combined fMRI and physiological monitoring and presented participants stimuli eliciting SET in a block-design study. A conjunction analysis identifying regions correlating with both the stimulus and heart rate, identified a set of brain regions including the pregenual ACC, orbitofrontal cortices, and the putamen. An analysis of the information pathways between these three areas using mediation analysis showed that they play independent roles in regulating the relationship between activity patterns related to a task's cognitive demand and those activity patterns related to ANS indicants such as heart rate. The non-artifactual nature of these findings was supported by the strong spatial localization of the effect, which stands in contrast to the more wide-spread correlates of physiological fluctuations. Based on the same data, other work (Wager, van Ast, et al., 2009) underlined the importance of prefrontal and subcortical systems as well: ventromedial prefrontal cortex and rostral ACC were shown to be involved in the relationship between brain function and heart rate, through the mediation of periaqueductal gray.

Another study (Gray, Rylander, Harrison, Wallin, & Critchley, 2009) explored the interrelations between ANS and neural activity using sensorial mechanisms. In that fMRI

study, several autonomic measures were acquired during the scan: electrocardiogram (ECG), HR and MAP data. Participants were instructed to wait for slight un-painful electrical shocks. These shocks were given either synchronously with or delayed with respect to the cardiac peaks (R-wave, dominant peak in EG cycle). This particular experimental setup was designed to examine whether stimulus processing would depend on a person's ongoing visceral state. Analyzing ANS indicants, they found that MAP increases when shocks were given synchronously with the ECG R-wave. BOLD signal was analyzing using regressors constructed by convolving the timing of synchronized and unsynchronized shocks with a hemodynamic response function. It was shown that the anterior insula, amygdala and brainstem showed different responses to synchronous and delayed shocks. Specifically, the left (and, statistically less significant, the right) anterior insula and the mid pons showed greater activity for synchronous shocks, and the right amygdale showed weaker activity. This suggests that there exist a set of brain regions that support the integration of somatosensory information with cardiovascular feedback to control autonomic arousal.

The spatial localization of the effects we previously described could be referred to the cortical back-ends of the interoceptive information pathways on the cortex [see (Critchley, 2005) for a useful schema]. The afferent information from the peripheral nervous system towards the central nervous system (Craig, 2002) passes through thalamic nuclei, and then towards several cortical areas: the insula, the anterior cyngulate cortex, the orbitofrontal cortices, and terminates in the right anterior insula (Craig, 2009; Critchley, 2005). The relationship between ANS and brain function emerge as a crucial factor in the BOLD fluctuations in these areas.

1.5 Summary

In this chapter we have presented two views of the functional role of ANS indices in the context of the fMRI BOLD analysis. On the one hand, these indicants are sometimes treated purely as physiological noise. When treated in this way, the aim of researchers is to remove the effect of these indices from the BOLD signal. As reviewed, a number of co-variables have been explored in this context, including the original cardiac and respiratory recordings, their harmonic expansions, under-sampled derivatives matched to the temporal resolution of the TR, and more complex derivatives capturing heart rate variance and respiration variance over time. Given the relation of such measures to purely nuisance variables such as head motion, cardiac-induced tissue motion and changes in CO₂ concentration, this approach is clearly justified.

However, the importance of work showing the interesting functional relation between ANS indices and both cortical and subcortical regions should be considered as well. From this perspective, brain regions that play a meaningful functional role in driving or monitoring ANS activity will show activation patterns that correlate with the BOLD signal. From this point of view, technical procedures that aim to remove the effect of ANS indices from the BOLD signal will also remove a meaningful variance component not only when studying processes associated with these regions, but when studying any process in which there might be a confound between cognitive processing and ANS states. The impact of “physiological cleaning” can therefore be particularly detrimental in various experimental contexts. These may include e.g., (a) manipulations of cognitive conflict that are associated with ANS activity, (b) studies that explicitly use emotional stimuli, (c) studies of decision making under uncertainty, (d) naturalistic viewing of engaging ecological stimuli such as movies that

induce fluctuations in ANS activity over time, and (e) studies evoking startle or surprise responses. Furthermore, as we have outlined, the study of the relation between the ANS and cortical function is an important topic being currently developed and describing the frequencies driving such connectivity, the pathways of information, and the way in which the link between ANS activity and cortical activity can be manipulated are central questions in this domain. (Clearly none of these questions can be answered if ANS indicants were removed from the signal.)

Addressing the role of the ACC and its function during rest is related to both these viewpoints. On the one hand, the ACC is a central node in the “Default Mode Network” (DMN) which is considered to be one of the dominant resting state networks (Buckner, Andrews-Hanna, & Schacter, 2008). Functional connectivity studies examining this region have treated ANS indicants as external factors that represent physiologically-confounding data. On the other hand, the literature we reviewed suggests that the ACC is one of the most important brain mediators between sympathetic information and brain activity, as was shown from studies of clinical and non-clinical populations. For this reason, it could be that the ACC mediates unique functions not shared with the rest of the DMN. For this reason, removing variance attributed by physiological measurements from this region could result in removing a unique component of variance, thus leading to an incomplete theoretical description of its role. For instance, it could be that the ACC is partly related to other regions mediating ANS functions, and this relation will not be found after the removal of ANS correlates.

On the basis of these considerations we suggest that researchers carefully evaluate whether their particular research question unambiguously justifies the removal of ANS

indicants from the signal. While it could be that in the near future methods will be available that will allow limiting the correction for ANS fluctuations to particular brain regions, to our knowledge such procedures have not been validated as of yet. For this reason, if the theoretical question being examined may have a functional link to ANS activity, it may be pertinent to analyze the data with or without the removal of ANS effects. In our opinion, increased synergy between research examining cortical systems mediating ANS function, and technical work aiming to identify ANS-BOLD correlates holds the potential to advance research in both domains.

In the following, we will address issues related to each of the two different perspectives we discussed in the first chapter, considering several issues related to each point of view.

In chapter 2 we will consider physiological quantities as generators of noise, and discuss conceptual flaws researchers have to face when dealing with data de-noising procedures. We first point out that it can be difficult to show that the procedure has achieved its stated aim, i.e. to remove only physiology - related components from the data. As a practical solution, we present a benchmark for assessing whether correction for physiological noise has achieved its stated aim, based on the principle of permutation testing.

In chapter 3, on the other hand, we will consider autonomic indicants derived from physiological time - series as meaningful components of the BOLD signal. There, we describe a FMRI experiment building on this, where the goal was to localize brain areas whose activity is directionally related to autonomic one, in a top - down modulation fashion.

In chapter 4 we recap the conclusions we found from the two approaches and we summarize the general contributions of our findings.

2 Physiology – related effects in the BOLD signal at rest at 4T

Summary

Blood-flow changes that are the source of the BOLD signal reflect a contribution of neurally-triggered fluctuations and those driven by physiological processes such as cardiac and respiration cycles. Physiological noise can have a potentially adverse impact on analysis of BOLD data. Beyond reducing the sensitivity of fMRI analysis, Physiological Noise (PN) can confound the interpretation of functional connectivity analyses, or studies of low frequency neural oscillations. For this reason, documenting cases where PN correction does not impact experimental findings is playing an important role in recent theoretical discussions: such cases suggest that connectivity is driven by functional activity rather than non-functional factors. A conceptual difficulty facing researchers interested in performing PN-correction is determining whether the correction procedure achieved its intended aim of removing PN components. If one cannot unequivocally show that the PN correction has achieved its stated aim, then any interpretation of its impact, particularly when referring to null results, is tenuous. We discuss these conceptual difficulties in the interpretation of data corrected for PN, and present a benchmark for assessing whether PN-correction has achieved its stated aim based on the principle of permutation testing.

2.1 Introduction

Fluctuations in the BOLD signal are driven by variations in cerebral blood flow (CBF) (Buxton, Uludağ, Dubowitz, & Liu, 2004). Changes in CBF are driven by neural activity but also by physiological mechanisms such as cardiac and respiration cycles. Recent work in cognitive neuroscience has begun paying attention to the relation between BOLD fluctuations and physiological mechanisms, and this interest is grounded in several reasons. From the perspective of signal analysis, BOLD fluctuations driven by autonomic processes constitute a nuisance component whose removal can improve signal quality and thereby increase the sensitivity or validity of statistical analysis conducted on humans or primates (Harvey et al., 2008; Lund, Madsen, Sidaros, Luo, & Nichols, 2006; Teichert, Grinband, Hirsch, & Ferrera, 2010). In particular, this procedure removes the potential impact of cardiac and respiratory effects that may be aliased into a low-frequency spectral domain as well as correct for non-independence of errors that is introduced due to these physiological components. This is a particular concern for higher field strengths where physiological modulations of intensity can become the dominant noise component (Triantafyllou et al., 2005). Another reason for the interest of physiological effects is grounded in the widespread interest in understanding the characteristics of inter-regional synchronization of BOLD time series ("functional connectivity"). To argue that a certain pattern of synchronization is driven by functional (i.e., neural factors) it is important to rule out the possibility that it is driven by physiological sources or alternatively, account for the potential impact of these sources on the BOLD signal. This is particularly the case when examining functional connectivity between brain regions in which BOLD activity correlates with physiological activity (Chang & Glover, 2009). The importance of using direct physiological regressors has been recently

demonstrated by showing that the use of proxy measures such as global mean to account for physiological processes introduces interpretive problems in and of itself (Murphy et al., 2009). In addition, quantifying and partialling out physiologically driven fluctuations is also important in the context of studies aiming to describe spontaneous neural activity associated with low- and ultra-low frequency fluctuations in the cerebral cortex (Chen, Tyler, & Baseler, 2003; Duff et al., 2008; Wu et al., 2008). Finally, understanding the relation between autonomic and BOLD activity is an important theoretical research topic in itself, in particular those studies aiming to understand how higher regions regulate autonomic process during behavior and rest (Critchley, 2005; Gianaros, Van Der Veen, & Jennings, 2004; Giardino, Friedman, & Dager, 2007) and (Gray, Minati, et al., 2009)[for a recent review]. This latter approach departs from prior ones in that it does not treat physiological information as a source of noise. For all these reasons, understanding the relation between BOLD activity and indices of physiological activity is becoming a shared interest in diverse fields of cognitive neuroscience.

Technically, accounting for the impact of physiological processes is relatively straightforward. Typically, cardiac and respiratory time series are acquired using dedicated equipment. From these time series several regressors are constructed. These may include the original cardiac and respiratory cycles and their temporal derivatives, but also respiration variance over time (Birn et al., 2006) or cardiac variance over time (van Buuren et al., 2009). The variance explained by these regressors is then subtracted from the BOLD time series, and the residuals of this procedure form a "physiological-noise corrected" PN-corrected version of the BOLD signal. PN-correction is therefore a simple technical procedure.

The question we are concerned with, here, is how can it be determined whether a given PN-correction procedure has indeed achieved its intended aim of removing physiologically-related variance from a BOLD data set? Demonstrating that PN-correction has achieved this aim is important in various research contexts such as those presented above. For instance, if one aims to rule out the possibility that BOLD synchronization in a certain network simply reflects physiological connectivity, one has to conduct PN-correction and in addition, independently demonstrate that variance related to physiological noise was indeed removed from the dataset. Only at that point is it most meaningful to discuss the features of the PN-corrected data. While there has been no detailed examination of this issue, it is clear that researchers examining the impact of PN-correction have been aware of its importance and addressed it in various manners. In what follows we briefly discuss the state of the art of these approaches and discuss their relative merits. We then present a detailed method for evaluating whether PN-correction has indeed worked effectively on either the individual or group level.

An examination of prior work utilizing PN-correction reveals that to date several procedures have been used, informally, to determine the effectiveness of PN-correction within different studies. Some (Bianciardi, van Gelderen, Duyn, Fukunaga, & de Zwart, 2009; Jones et al., 2008) have evaluated the impact of PN-correction on a group level by quantifying its impact in terms of the difference between a time-series' variance prior to and after correction (e.g., $1 - [\sigma(\text{post}) / \sigma(\text{pre})]$). The rationale behind this method is clear: if the procedure does not induce a change in the time series variance, then it is likely that PN-correction has not achieved its aim. However, when reductions in variance are identified using this procedure their interpretation is not straightforward: variance reductions can be

spuriously caused by the regression procedure underlying PN-correction. This method often relies on applying a large number of regressors which can lead to a reduction in variance by chance alone. Thus, documenting that PN-correction reduces variance by a “reasonable” proportion, heuristically defined, might be insufficient for concluding that the procedure was performed correctly.

A different method for assessing validity is based on an individual-level analysis. Here, for each participant the overall fit of the PN-correction regression model and the BOLD time series is derived (and assessed via an F-test). Then, the number of voxels for which the fit exceeds chance is calculated (Chang & Glover, 2009). If the PN-correction model fit is reliable for more voxels than be expected by chance, this is taken to indicate that PN-correction has achieved its aim. For instance, if 10% of the voxels show a fit with the PN-correction regression at an alpha threshold of $p < .05$, this would indicate that there are twice as many voxels as would be expected by chance. This method, however, brings its own complications. Due to the serial autocorrelation with both the BOLD and physiological time series, the degrees of freedom used in such tests are inflated relative to the effective degrees of freedom within the time series. This leads to an overestimation of the F value and its reliability. In other words, the use of such tests could make the PN estimators to appear to be more accurate than they really are.

In order to take in account these issues and according to the question we stated before we thought about a method that, in a self-consistent way, assess the validity of the PN-correction procedure, and we applied it to a resting-state dataset. The core of the method mainly relies in building sampling distributions of quantities that reflect the quality of the data to compare distributions obtained with real cleaning procedures.

Here we want to illustrate this method and show how it resulted straightforward in answering validly the scientific question we stated before, avoiding all the issues we described: first, the cleaned data are compared not only with the raw set. We instead build, for each participant, a sampling distribution of a variance-related quantity that is the signal to noise ratio of the datasets cleaned with other participants' physiological data. We also describe the simulations we made to extract trends in the improvements as a function of several parameters of the regression, in order to have amounts of variance that are not due to any specific effect in particular, to compare with real regressions. Furthermore we illustrate several issues about the autocorrelation of the time-series we use to clean data, and we relate it to the results we obtain with the cleaning procedure. With this description we claim to show that new elements need to be introduced in the debate about physiology and fMRI, and we propose a self-consistent method to clarify when to accept or to reject datasets that underwent a cleaning procedure.

2.2 Methods

2.2.1 Imaging

30 adult volunteers (15 female, aged 27.8 ± 8.5 years) participated in the study. All volunteers provided informed consent to their participation in the study which was approved by the institutional ethics committee of the University of Trento. Images were acquired with a 4T Bruker scanner. Functional scans were acquired for each subject using the following parameters: TR/TE=1500/34 ms (172.5s overall); matrix=64x64; 25 parallel to AC-PC slices; voxel size: 4x4x(4+0.8 spacing) mm. Structural scans were acquired with a 3D T1-weighted MPRAGE sequence: TR/TE=2700/4 ms, isotropic voxel size = 1mm, matrix=256x224; 176 sagittal slices. 15 participants underwent a functional scan session of 115 scans (172.5 s).

The other 15 underwent a functional scan session of 215 scans (322.5 s), that was split in two datasets. We ended up with 45 resting-state datasets of 100 scans each. In the resting-state sessions participants were instructed to passively fixate on a white crosshair placed in the center of the screen, to remain awake and attempt not to think about any topic in particular. Cardiac and respiratory processes were monitored using a photoplethysmograph placed on a finger of the left hand and a pneumatic belt strapped around the upper abdomen [Bruker] respectively. Cardiac and respiratory data were both sampled at 50 Hz and saved to a file while acquiring MRI images.

2.2.2 fMRI pre-processing

All analyses were conducted in the original subject space and voxel dimensions. No registration to common template was performed. Analyses were performed using the AFNI neuroimaging tools (Cox, 1996). The target voxels of interest in all analyses were those located within the intracranial region. To identify these regions the structural image was aligned to the functional EPI image using an automated algorithm (Saad et al., 2009). Alignment was verified manually and adjusted as needed. A skull-stripping procedure was then applied to create a mask marking only those voxels falling within the intracranial mask. The time series of these voxels were used in all subsequent analyses. No temporal smoothing or temporal filtering was employed, and no spatial smoothing was performed on the data. First 15 scans were removed from the initial sets of images to allow the MR signal amplitude to stabilize: we worked with sets of 100 scans each. Then functional data were pre-processed using AFNI suite. Voxelwise temporal despiking (removal of the outliers with respect the temporal median, replaced with values interpolated with the nearest-neighbors') and a

regression of global trends subtracting effects explained by 0th to 3rd order polynomials were performed.

2.2.3 Construction of the regressors and quantification procedure

The core of the data manipulation was the removal of physiology-related effects with a slice-based regression: this procedure consisted in construction of physiological regressors and the subtraction of their effect from the BOLD signal. The construction of the slice-based set of regressors was performed as follows: we used AFNI-Matlab code (RetroTS.m, Ziad Saad) to construct 13 slice-based regressors from the 2 initial cardiac and respiration time series. For each participant this generated a 13(regressors) x 25(slices) matrix. The 13 regressors generated from the cardiac and respiration time series included: 5 Respiration Volume per Time (RVT), 4 respiration and 4 cardiac time-series, with a sampling frequency that matches the fMRI data one (Birn et al., 2006; Glover et al., 2000). After constructing these regressors we subtracted their effects from the BOLD time series of each voxel using a regression procedure (Glover et al., 2000). This generated a functional dataset where the variance explained by the set of regressors was removed. Henceforth, we refer to this regression procedure as PN-correction. When a participant's own physiological data were used for constructing regressors to be applied to that participant's BOLD time series we refer to the result of correction as the Real-PC dataset. To construct a sampling distribution of the impact that PN-correction can generate by chance we used a permutation method as detailed in the following section.

2.2.4 Permutation method

To obtain a reference sampling distribution against which to compare the Real-PC dataset of each participant we constructed surrogate datasets via permutations. A permutation

consisted of using another participant's physiological time series as inputs to the PN-correction procedure. That is, given N sessions, $N-1$ surrogate datasets were generated via permutation for each participant. The advantage of this permutation procedure is that it maintains the temporal characteristics of physiological time series [which often have complicated spectral properties (Hu, Lee, Gao, White, & Crosson, 2008), including serial autocorrelation]. In our particular study we obtain 44 surrogate datasets for each dataset, and we refer to these as Perm-PC data sets.

To evaluate the effectiveness of the PN-cleaning procedures we chose the temporal signal to noise ratio (tSNR) as a standard voxelwise quantity for evaluating the Real-PC and Perm-PC data sets. tSNR is defined as the ratio between the temporal mean and the temporal standard deviation of a time-series. This quantity is an index of the quality of the data; the higher the tSNR, the better the quality of the time-series. We evaluated the tSNR measure for the three types of datasets in the current study: (1) the raw data, that was elaborated only with the procedures described in the “fMRI pre-processing” section but that did not include PN-correction of any form; (2) The 44 Real-PC datasets in which PN-correction was performed by applying a participant's physiological regressors against their own BOLD data; (3) The 44x45 (1980) Perm-PC datasets, where we applied the PN-correction on BOLD datasets by using other participants' physiological data.

Given the specific properties of the physiological acquisitions we wanted to assess whether the autocorrelation influenced significantly the result. We took original physiological acquisitions and scrambled them, to obtain two random time-series. Then we used them as physiological data, building set of 13 regressors for each dataset. We then went on in a way similar to that we described before, ending up with 45 set of 45 simulations

(total=2025 datasets). We performed quality tests on these dataset, and made comparisons with those coming from the real one.

We finally constructed a simulation to assess the basic effects of a regression procedure on completely random data, varying parameters such as time-series length and regressor number.

2.3 Results

To assess whether a procedure worked in cleaning from the variance related to physiology we followed several steps to characterize and attack all the issues of the problem.

2.3.1 Simulations

First, we observed that the computations behind all the procedure involve several kinds of systematic confounding factors, and in order to understand them we built a simulation. We constructed sets of 1000 random time-series and 13 random regressors. Time-series length (100 time-points) was the same of our real data-sets. The design matrix included an intercept that take in account the mean value of the time-series. We evaluated residuals from the regression and we obtained, running the simulation 100 times, (7.41 ± 0.08) % of tSNR change, and (13.11 ± 0.13) % of variance saved. This suggests that some amounts of change in the variance could be due to a plain statistical effect. We studied also their trend as a function of time-series' length and we noticed that the longer the data-set, the lower the amount of variance removed by random regressors. We finally manipulated the number of regressors, obtaining a gradual increase of the amount of variance explained, as the number of regressors increased. This indicates that saturating the model with more and more regressors, in order to obtain a better description of the physiology effect, could lead to an

artifactual increase of the quality of the data. All these simulation results could help in pointing out when there is a random effect, and when the cleaning procedure worked in explaining the correct variance.

2.3.2 Permutation tests: comparison between cleaning procedures

We report the overall quality of a dataset as the mean value of the temporal signal to noise ratio, evaluated voxel-wise. The particular definition of this quantity allowed us to deal with an important feature of the signal that is the spread of the fluctuations around the mean. We show the results in table 2.1: a way to assess whether the cleaning procedure worked is to compare the tSNR of the data cleaned from subject's own physiological data (Real-PC, 9th column) with the raw time-series quality (2nd one). We obtained that for all the subjects the cleaning procedure improved the tSNR in a statistically significant ($p < 10^{-6}$, corrected for multiple comparisons) way.

However, if we insert as a term of comparison the datasets obtained from the permutation framework, 8 out of 45 dataset did not pass at least one of our control tests. We first excluded datasets where the Real-PC tSNR was less than the best permutation: in this way we claimed to mark datasets where the procedure did not improved the quality of the data better than a procedure fed with regressors unrelated to subject's specific fluctuations. 7 datasets did not pass this test. We then performed a test to assess significant differences between Real-PC datasets that previously worked and best permutations. This led to the exclusion of 1 more dataset.

ID	Mean Raw tSNR	SD Raw tSNR	Mean Scrambled tSNR	SD Scrambled tSNR	Mean Permutation tSNR	SD Permutation tSNR	Mean Real tSNR	SD Real tSNR	
1	108.92	55.95	114.28	60.33	127.28	64.62	134.06	66.50	
2	92.90	50.86	96.92	53.87	108.44	58.59	115.73	60.57	
3	52.54	35.73	84.62	47.28	101.78	53.20	103.20	53.30	
4	74.66	44.33	81.36	47.76	95.95	53.20	94.08	52.91	X
5	89.79	49.63	102.22	54.44	110.92	57.28	118.34	59.50	
6	95.46	48.03	99.50	51.53	107.36	54.11	128.62	57.68	
7	91.01	46.76	95.04	50.11	102.92	52.65	123.45	57.20	
8	106.29	53.86	114.45	57.94	125.22	62.00	136.05	64.70	
9	78.95	44.84	74.11	44.85	97.38	51.67	96.65	53.30	X
10	114.09	54.63	119.78	58.61	134.49	62.84	135.57	63.16	
11	91.49	48.86	97.60	52.61	108.46	56.06	114.08	57.57	
12	124.91	56.83	128.51	60.71	147.52	65.60	155.82	66.36	
13	81.31	44.89	78.25	45.87	98.82	52.53	102.45	53.13	
14	46.16	31.95	46.42	32.75	59.49	38.66	58.15	39.01	X
15	57.42	36.31	60.25	38.33	69.10	42.34	74.92	44.60	
16	113.65	55.36	117.88	58.41	130.05	62.57	136.81	64.24	
17	93.68	53.16	102.40	55.73	112.78	59.90	118.51	61.85	
18	115.74	53.76	118.46	57.61	132.79	61.54	152.96	62.76	
19	110.51	52.13	113.84	55.89	130.91	60.10	143.68	60.67	
20	104.78	51.25	108.35	54.51	118.17	57.69	137.23	59.29	
21	83.35	44.26	91.12	47.56	101.89	52.21	109.45	53.13	
22	68.33	41.76	67.01	43.25	84.22	48.66	85.15	49.77	
23	75.35	44.28	76.05	45.50	94.92	52.44	96.50	52.82	X
24	83.16	45.46	83.77	46.53	101.84	53.14	101.09	52.85	X
25	92.29	52.83	96.41	56.62	107.97	60.25	118.89	63.12	
26	117.71	57.04	121.90	60.78	134.31	64.66	132.64	64.42	X
27	119.21	57.33	121.52	61.19	136.46	65.69	133.48	65.05	X
28	84.27	44.36	89.10	47.13	101.37	51.63	107.84	52.48	
29	92.95	46.97	98.34	50.32	105.83	53.41	116.30	55.08	
30	86.34	49.93	92.68	53.85	99.82	57.09	122.37	61.92	
31	118.52	56.24	124.89	60.96	135.59	64.07	148.07	66.21	
32	106.69	53.95	112.83	58.00	124.49	61.79	132.54	64.12	
33	48.39	29.11	50.37	29.26	71.09	37.37	71.89	37.76	
34	86.37	46.17	88.57	48.29	102.62	53.04	103.32	53.57	
35	85.71	46.01	88.57	49.14	101.32	52.53	101.96	53.66	
36	73.67	40.43	78.05	42.57	87.98	46.57	96.34	48.56	
37	76.87	43.75	85.13	47.07	91.83	49.23	102.80	52.71	
38	112.34	56.65	119.26	60.77	128.35	64.67	144.48	66.79	
39	111.70	56.00	117.99	60.57	128.14	64.13	140.30	65.44	
40	80.02	46.74	88.86	50.74	96.98	53.80	107.50	56.19	
41	75.86	44.08	83.16	47.55	92.34	51.33	103.06	53.85	
42	88.38	47.47	88.83	48.16	111.58	58.48	104.31	54.49	X
43	95.28	47.57	100.87	51.27	107.79	53.31	122.29	55.88	
44	84.83	44.71	89.33	47.93	98.22	50.84	107.06	52.87	
45	107.62	55.98	110.11	58.36	124.18	64.18	136.77	64.86	

Table 2.1 – Comparison of signal – to – noise ratio evaluations for each of the datasets. The last column on the right reports a X when mean tSNR calculated using participants’ own physiological data is not significantly greater than the one evaluated using the permutations.

A practical way to graphically visualize when the procedure of cleaning worked is to plot the cumulative distributions of the tSNR of the brain, comparing curves from different elaborations. We show these results in figure 2.1, comparing a participant where the procedure worked (A), and one where the procedure failed (B). Each of the two plots refers to one dataset. On the x-axis there is the tSNR, on the y-axis we plot cumulative distributions of tSNRs evaluated in 4 conditions: Raw (red), Scr-PC (black), Perm-PC (blue) and Real-PC (green). The plots reflect what we reported in the table (rows: 21 and 27). Starting from the Raw curve, where the data quality is lower, and moving towards the right (higher data quality according to the tSNR metric), we notice that in participants where the procedure worked (1a) we encounter at first the black curve, then the green set of curves and finally, far on the right with respect the red one, the blue curve. In participants where the procedure did not work (1b) we cannot distinguish whether the blue curve is on the right with respect the green set.

Black curve has several features that help in understanding what happens to the time-series undergoing a cleaning procedure. First, the scrambled time-series present low autocorrelation, so their improvement in the tSNR of $5.35 \% \pm 2.90 \%$ is comparable to the one obtained from a completely random simulation described previously. We plotted one curve coming from all the datasets given that there is no distinction between scrambled regressors. Furthermore, the distance between the black curve and those on left and right is an index of how much variance a regression accounts by itself: we noticed that the overall improvement from raw tSNR to real one is arranged in a way that reflects the properties of the regressors used. The elaboration fed with scrambled regressors accounted for about 20%.

Cumulative distribution of temporal signal-to-noise ratio in two cases

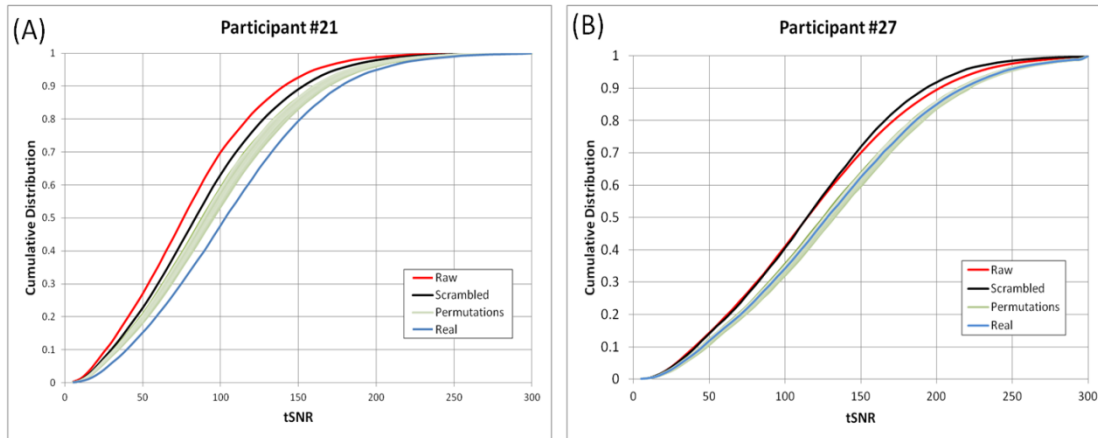


Figure 2.1- Cumulative distributions of tSNRs evaluated in 4 conditions: Raw (red), Scr-PC (black), Perm-PC (blue) and Real-PC (green). The plots reflect what we reported in the table (rows: 21 and 27). Starting from the Raw curve, where the data quality is lower, and moving towards the right (higher data quality according to the tSNR metric), we notice that in participants where the procedure worked (A) we encounter at first the black curve, then the green set of curves and finally, far on the right with respect the red one, the blue curve. In participants where the procedure did not work (B) we cannot distinguish whether the blue curve is on the right with respect the green set.

Then the most consistent part of the increasing in tSNR belongs to the permutation set, accounting for about 54% of the improvement. According to this arrangement, the amount explained only by subject’s own physiological fluctuations is about 26%. This suggests that most of the improvement of the quality of the data is given by specific properties of the regressors, and the effect of specific fluctuations is not actually the one that increases the tSNR. It results crucial when deciding where the procedure worked, if compared with bulk improvements given by elaborations fed with other regressors.

2.3.3 Permutation tests: changes in tSNR

To determine whether PN-correction induces a non-uniform improvement in tSNR as function of initial tSNR, we evaluated the relation between a voxel’s original tSNR (pre-cleaning) and the percentage-improvement in tSNR introduced by the cleaning procedure. tSNR improvement was defined as $-100 \cdot (\text{tSNR}_{\text{pre}} - \text{tSNR}_{\text{post}}) / \text{tSNR}_{\text{post}}$. Figure 2.2 shows

the relation between initial tSNR ($tSNR_{pre}$) and that resulting from PN-correction ($tSNR_{post}$) for three variants of PN-correction: one based on random regressors (black bars - Scrambled-PC), one based on permutations (green bars - Perm-PC) and one based on participants own PN data (blue bars - Real-PC). Bars summarize blocks of 25 unity of tSNRs: we therefore sampled 12 blocks of tSNR from 0 to 300.

Changes in standard deviation of the time-series in two cases

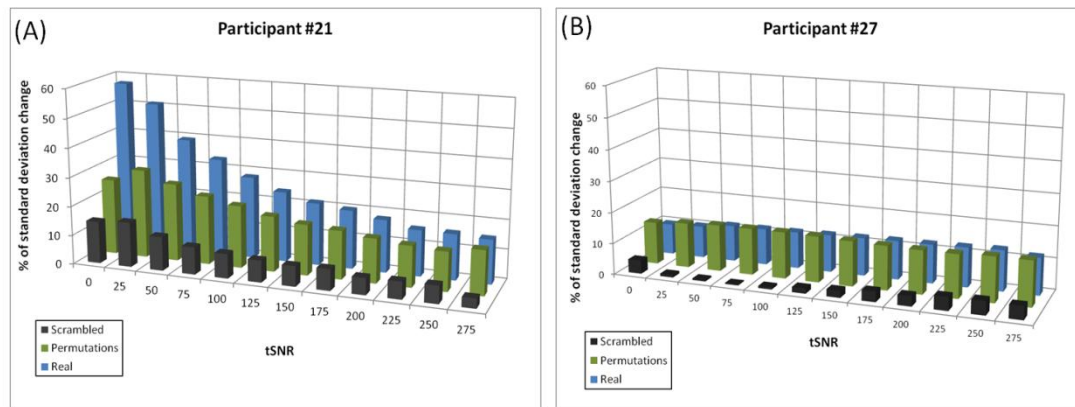


Figure 2.2 - Relation between initial tSNR ($tSNR_{pre}$) and that resulting from PN-correction ($tSNR_{post}$) for three variants of PN-correction: one based on random regressors (black bars - Scrambled-PC), one based on permutations (green bars - Perm-PC) and one based on participants own PN data (blue bars - Real-PC). Bars are shown for one subject where the cleaning procedure turned out to be valid (A) and for one where the procedure did not work (B).

The impact of PN-correction using participants' own data is evident in a non-uniform pattern where the impact of the cleaning procedure is negatively correlated with a voxel's initial tSNR. That is, the lower a voxel's initial tSNR, the strongest the impact of PN cleaning.

This trend was assessed, for each dataset, by a linear regression of the tSNR changes against the initial tSNR. The negative linear trend was reliable ($P < 0.05$ corrected for multiple comparisons) in 97% of the datasets. This result was also significant on the group level ($t(44)=-4.9$, $P=10^{-5}$)

A weak negative correlation linear trend was also introduced by the permutation procedure. Negative linear trend was reliable in 95% of datasets and at the group level ($t(44)=-2.8$, $P=0.0087$) too.

A direct contrast revealed that the slope of the linear trend was always (100% of the datasets) significantly ($P<0.05$ corr.) stronger for the Real-PC than Perm-PN datasets. This holds ($t(44)=6.14$, $P=2*10^{-7}$) at the group level too.

To compare the magnitude of the slopes we calculated the mean weighted with the standard deviations, in order to consider the spread of the results for each subject. We ended up with a mean slope $(8.23 \pm 0.04)*10^{-2}$, for the Real-PC dataset, that is about 3 times the Perm-PC one $(2.72 \pm 0.02)*10^{-2}$.

We considered the results described in pictures 1 and 2 complementary and mandatory to assess if the procedure to clean for the correct variance worked or not. From this perspective, plots of tSNR change have two important issues: first they describe a general property of the cleaning procedure, and second they are useful in particular situations, for example when it is difficult to disentangle the results from Perm-PC and Real-PC (for several reasons: e.g. high cross-correlation between one particular permutation and the current physiological data), to take a look to the changes with respect to the raw values could be a way to force the assessment of the goodness of the cleaning procedure, showing a difference in the decreasing trends.

2.3.4 Autocorrelation of the regressors

The core of the method we described is the comparison between a result (Real-PC) obtained combining fMRI data with regressors coming from the same experiment, and a

sampling distribution (our empirical chance level) thought as something that comes from experiments as well. We studied the possibility to compare Real-PC results with some obtained from completely random time-series: in this way we avoid the temporal autocorrelation that characterize both the physiological time-series and the fMRI data. In figure 2.3 we show a plot for each of the 13 RETROICOR regressors averaged over all the subjects we used.

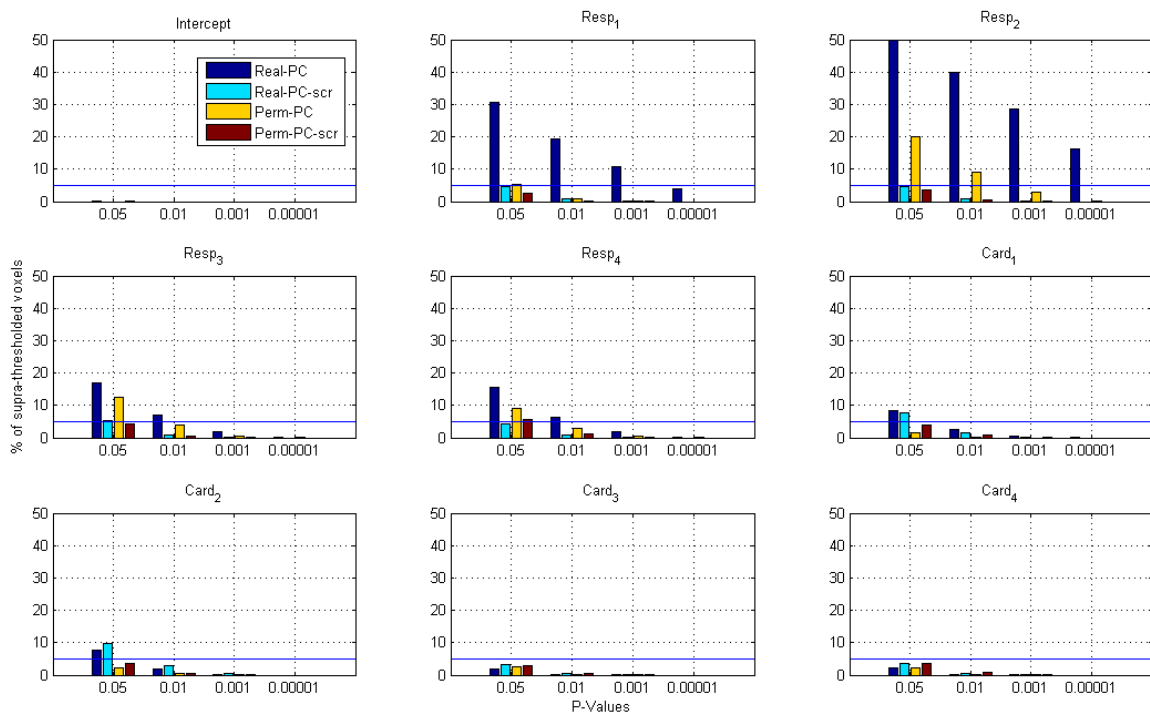


Figure 2.3 – Percentage of supra – thresholded voxels for each RETROICOR regressor, comparing procedures fed with participants’ own physiological data and permutation one. We report also bars referring to the same regressors scrambled in time. Overall, regressors referring to respiration seem to be the relevant ones. Moreover, as far as permutation bars are greater than scrambled – real ones, autocorrelation in the time – series seem to be a significant feature in driving the method.

For each of them we report 4 bars for 4 p-values. Bars are the percentage of significant voxels that pass the p-value. Dark blue bars refer to Real-PC procedure; light blue bars refer to the Perm-PC procedure where one subject fMRI data are cleaned from the activity explained by other subjects' physiological data. Yellow and red bars refer to regressors built

on scrambled physiological time-series. We could extract two different pieces of information from this plot: first, given that the dark blue bars are coming from the same (Real-PC) regression, we could estimate what are the most relevant regressors involved in the cleaning procedure. It seems that respiration-related (RVTs, Resps) quantities explain more variance than the cardiac. Second, the autocorrelation seems to be an important feature of the almost-random regressors we built to compare the real regression with. Dark blue bars are always bigger than the Perm-PC (as we observed in the previous results). This means that, as an overall result coming from the unscrambled regressors, the Real-PC procedure works better than the Perm-PC. Moreover, Perm-PC bars are always bigger than the chance level indicated by the respective p-value. This means that almost-random regressors explain an amount of variance that is bigger than the random one. We observed that there is some cross-correlation between regressors that could result, in some datasets, in variance explanation from physiological data of other participants. Furthermore Real-PC bars are bigger than the yellow and red ones as well: the Real-PC-unscr procedure works better than the two regressions fed with scrambled regressors. This could suggest that scrambled, completely-random time-series could be used to assess whether the cleaning procedure worked or not. A comparison between the amount of regressors' autocorrelation and the number of supra-thresholded voxels for each unscrambled regressor shows that there is not a direct relationship on a single-regressor level.

In figure 2.4 we show the amount of voxel saved when performing the cleaning procedure fed with regressors built from the scrambled physiological time-series: we considered fMRI dataset where the variance explained by these random regressors was removed, and we evaluated how many voxels experienced this effect as non-random.

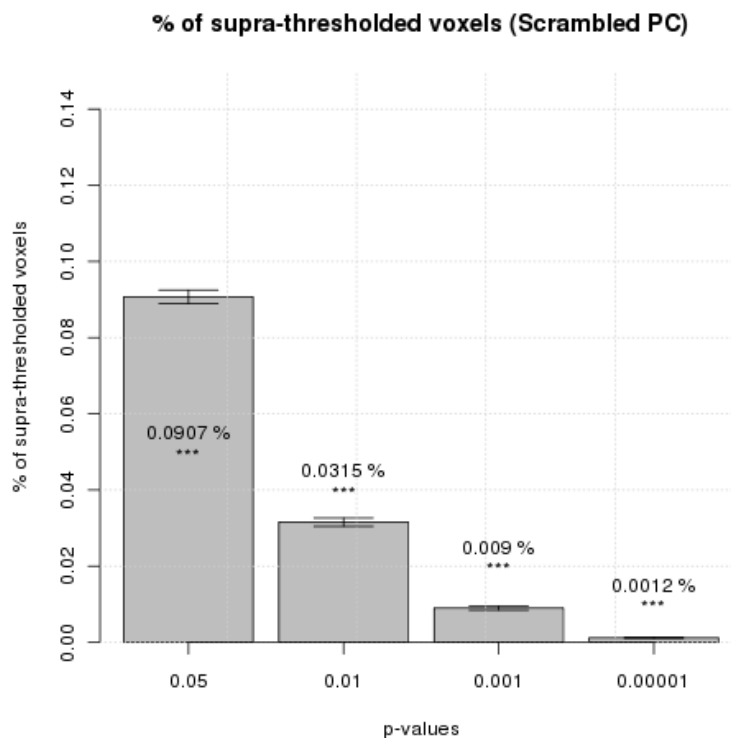


Figure 2.4 – Percentage of supra – thresholded voxels obtained after performing a correction using random (scrambled) time – series.

We performed this sort of simulation for each subject, using all the time series we had. We finally obtained 45 datasets cleaned from 45 set of regressor each, resulting in a total of 2025 results. We extracted the percentage of voxels where the procedure was shown significant obtaining a collection of percentages according to 4 p-values, from 0.05 towards 10^{-5} . We report the average of the percentages on the bars, the stars indicating the significance of the bar from the current chance level. We observed that the amount of voxel saved is always significantly different from the actual randomness. This suggests that using regressors that come from completely random physiological time-series lead to results that are artifacted by the way in which the regressors are built.

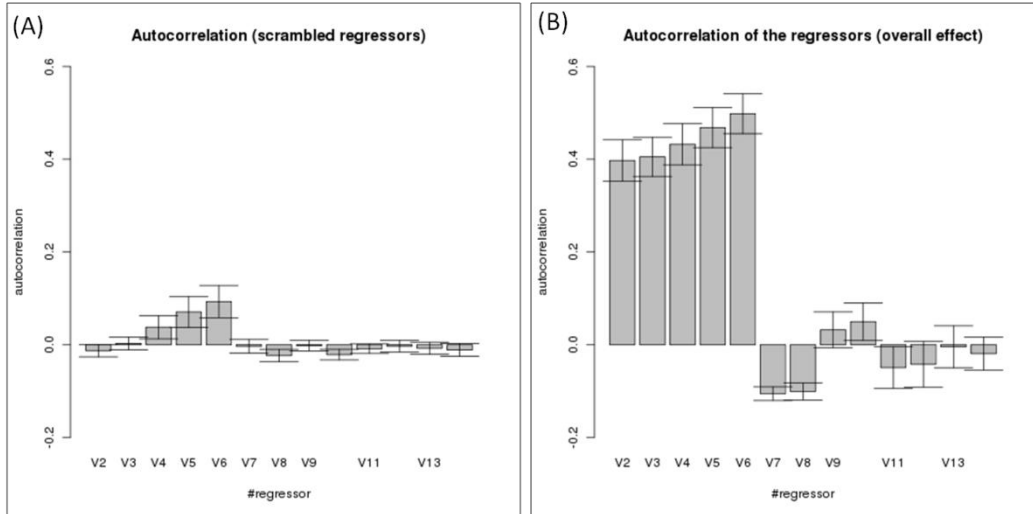


Figure 2.5 – Comparison of the amount of autocorrelation in (A) scrambled and (B) original regressors. The latter set is also used to build the sampling distributions in the permutation procedure.

Specifically, in figure 2.5 we compare the autocorrelation levels between our two non-real sets of regressors. In panel (A) we plot the amount of autocorrelation for scrambled regressors, in (B) for the permutation set. We notice that the amount of autocorrelation, in the latter, is about one order of magnitude greater with respect that in the scrambled set of regressors. This suggests that the goodness of the cleaning procedure, in general, could be due also to this signal property.

2.4 Discussion

In the present study, a method to assess the validity of a PN-correction procedure was presented. This originated from several issues we noticed when looking for a correct characterization of the variance in the BOLD signal, in order to leave only what would be relevant for further analysis. We found instead that the topic is intrinsically full of interesting aspects, and worth to have a thorough discussion, even without a combination with cognitive issues.

Here we suggested that a plain comparison of standard deviation values before and after the correction, could potentially generate misses, and result in a liberal assessment of the procedure. Simulations and Table 2.1 help in going through this point.

We reproduced some features of our dataset, using completely random time-series and regressors used in the analysis (100 time points and 13 regressors, for each voxel). We ended up with a significant percentage of standard deviation decrease (about 7%). This result is a function of the number of regressors and, above all, of the length of the data-set. It represents a sort of ground-value to compare all the other with, given that here reside all the effects that don't depend on the quality of the data. We suggest that building a simple procedure like this, one for a whole sample of datasets, could be an important comparison to quickly assess whether the PN-correction produced a result different from the chance.

We investigated the concept of chance itself, introducing the sampling distribution built cleaning one subject's fMRI data with other participants' physiology data.

A comparison of the 2nd and 9th column in **Table 2.1** returned us the fact that all the Real-PC mean values for the tSNR were significantly different from the raw one. In this case, comparing only these two quantities, one could conclude that the quality of all the dataset was increased and the procedure worked. Using the simulations we had a first control to exclude those datasets where the change in standard deviation did not result different from the chance. The random regressors don't take in account several features that pertain to physiological data, like the serial autocorrelation.

Thus the sampling distribution built with the permutation of physiological data seems to be a significant way to compare the Real-PC correction with. Conceptually, we built a

baseline with several features, that resulted itself different from zero. Data cleaned with other participants' physiological data have better quality, in the sense of bigger tSNR, with respect the raw dataset: they represent a significant manipulation of the data. But the actual goal of the PN-correction is to remove variance that is related to physiological fluctuations. So the comparison of real-PC with perm-PC set is, in this framework, the contrast we consider the most valid.

This led, however, to the rejection of as many as 8 datasets, out of 45: more than one sixth of the total. We considered this issue exploring why several datasets coming from the permutations were so close to the real-PC data.

In figure 2.3 we plot the amount of variance explained by each of the 13 regressors. The information carried out by these plots resides in the fact that yellow and dark blue bars are, overall, significantly greater than the other bars. The main difference between these bars and the other two is in the autocorrelation. We open up this point in figure 2.5. There we show an important comparison between the random and the real regressors. The former set, shown in (A), has no significant autocorrelation, while the real regressors, whose autocorrelation is in (B), show a significant pattern of this specific feature, which could be driving the permutation procedure.

In figure 2.2 we report trends of tSNR changes as a function of initial quality of the data. We discussed that this result is complementary of the one showed in figure 2.1. The trends reflect the fact that the physiological data exhibit common spectral properties for all the participants. There is still a difference in the strength of the slope, the real-PC one resulting overall stronger than the perm-PC's. This could work as a sort of second level check of the

results obtained comparing tSNRs. If the real-PC signal quality does not come out significantly different from the perm-PC one, but the trends for tSNR do, then the procedure could be still judged as working.

We cited, introducing the work, that beyond the simple comparison of standard deviation amounts, there are other ways to show the effects of the elaborations. Counting the supra-thresholded voxels could be a straightforward way to deal with this. Nevertheless we show in figure 2.4 that random regressors could lead to significant amounts of voxels saved.

2.5 Conclusions

The considerations we made and the results we showed could drive to the fact that the problem is not the way in which we look at fMRI data cleaned from physiological effects. We instead think that the real issue is to find a significant chance level to compare the results with. We tried using completely random regressors, but the autocorrelation feature seems to be dramatic when explaining BOLD signal fluctuations. The choice of the permutations framework seems to be proper for several motivations. First, other participants' physiological data are a ready-to-use information that comes from the experimental sessions themselves, with no need of further acquiring. Then, they are on one hand random, because it is unlikely that two subjects, even if they share heart rate and respiration frequencies, report the same trend of these fluctuations in time. On the other hand they still exhibit all the features that these data should have: e.g. autocorrelation and possible artifacts, given that they are taken from the same hardware.

We therefore think that the sampling distribution built from experimental data is an issue that gave us a reliable chance level to compare real data with, maximizing the effects we listed as potentially responsible to loss of validity.

The distance between real-PC result and perm-PC one seems to be a straightforward choice to characterize those datasets where the changes in variance are likely to be related to the physiological time-series. Using this distribution we were able to show, in a reliable way, where those changes could instead reflect computation or other intrinsic features of BOLD signal.

3 Directional relationships between BOLD activity and autonomic nervous system fluctuations revealed by fast fMRI acquisition

Summary

Identifying neural systems where BOLD fluctuations are related to near - future autonomic activity (AA) fluctuations is an aim that could potentially contribute to three general questions: first, are there brain areas where cognitive activity is followed by systematic changes in autonomic activity. Second, to the extent that such are found, do they perform similar functions during task and rest states. Finally, are the brain regions showing this relationship brain regions that are considered active by task, or do they also include brain areas deactivated by the task. To answer these questions in a valid way, we constructed a fast – repetition time (fast-TR) functional magnetic resonance imaging experimental setup, combined with physiological acquisitions, and also included a block-design experiment with a typical TR to quantify whole brain activation and deactivation during the task.

Participants underwent two resting state scans (one under a fast-TR setup and one under a slow-TR setup) and also underwent several task scans under slow- and fast-TR scanning conditions. We analyzed the data using a lagged-correlation method between the BOLD signal and autonomic activity related timeseries. Using this method we found that anterior

cingulate and insular cortices have important roles in BOLD – autonomic activity exchanges, but their relation to AA activity is quite different. We observed two areas considered important hubs for the brain activity at rest, such as posterior and anterior cingulate cortices, playing different roles with respect to the autonomic activity. Finally, we cross-referenced the regions identified as driving ANS activity with the set of regions active and deactive during the task and found that BOLD -> ANS relations are found in both task-active and task-deactive regions, a finding that has important implications for the type of processing occurring in deactive regions.

3.1 Introduction

3.1.1 Immediate background

Theories of neurovisceral integrations, i.e. those attempting to explain connections between brain and peripheral organs including the heart, assume the existence of general brain systems whose function is to integrate external and internal information in order to monitor external environment (especially safety and threats) and prepare actions (Thayer, Ahs, Fredrikson, Sollers, & Wager, 2012; Thayer & Lane, 2009). This neurovisceral system could be divided in three different and almost hierarchical modules. First, a cortical subsystem including prefrontal cortex and, subsequently, cingulate cortex and insula. Then, an intermediate subsystem including the nuclear structure of the amygdala, which communicates with the second module, involving subcortical structures and third the autonomic module, controlled by non-coupled (G G Berntson, Cacioppo, & Quigley, 1991) sympathetic activation and parasympathetic (vagal) suppression. The cortical module, which is the relevant one for the purposes of this work, includes areas like the cingulate cortex, also known to show non-random BOLD fluctuations at rest. The autonomic module, on the other

hand, includes vagal and sympathetic modulation of heart rate: these could be investigated using heart rate variability as a simple and valid indicant. This neurovisceral integration hypothesis draws on and expands an important body of work done by Critchley and colleagues (Critchley et al., 2003), where it was shown, using a comparison between healthy participants and patients with lesion in anterior cingulate cortex (ACC), how activity in several cortical areas varies with changes in arousal, and how autonomic indicants carry behavioral information too.

The problem of characterizing BOLD fluctuations at rest (Raichle et al., 2001) is one of the most debated topics since early 2000s. There are three interesting issues related to resting – state, for the purposes of the present work: first, the anterior cingulate cortex is a central node among the set of areas known to show correlated, non – random fluctuations at rest. It is an important part of the so-called ‘default mode network’. The fact that the ACC might be involved in monitoring or triggering ANS activity is highly pertinent to interpreting the nature of its activity during rest, because such non-random resting-state patterns could be related in an artifactual way to BOLD fluctuations (Birn, 2012). Alternatively, it could be that part of the information processing carried out in the ACC during rest has to do with triggering downstream autonomic nervous system (ANS) activity. As described in chapter 1, the relation between autonomic indicants and ACC and so-called default mode network (DMN) activity has been noted before, and typically the time series of ANS indicants is regressed from the ACC signal recorded with the MR instrument. Our investigation can contribute to understanding ACC functions during rest by evaluating whether it fluctuates synchronously with ANS oscillations (which would be consistent with an artifactual impact of ANS → BOLD), or whether it precedes ANS fluctuations. Finally, beyond identifying the

BOLD ↔ ANS relation in ACC (and other regions) during rest, we also aimed to quantify their relation during task execution. The ACC and posterior cingulate cortex (PCC) are regions known to be deactivated during task performance; that is, BOLD signal in these regions is lower during task performance than during attentive rest. Understanding whether these regions show BOLD → ANS relations during a task in which they are deactivated would constitute an important contribution to our understanding of the processes that take place during task-induced deactivation.

3.1.2 Design principles

In what follows we present the motivations for the different components of the experimental protocol.

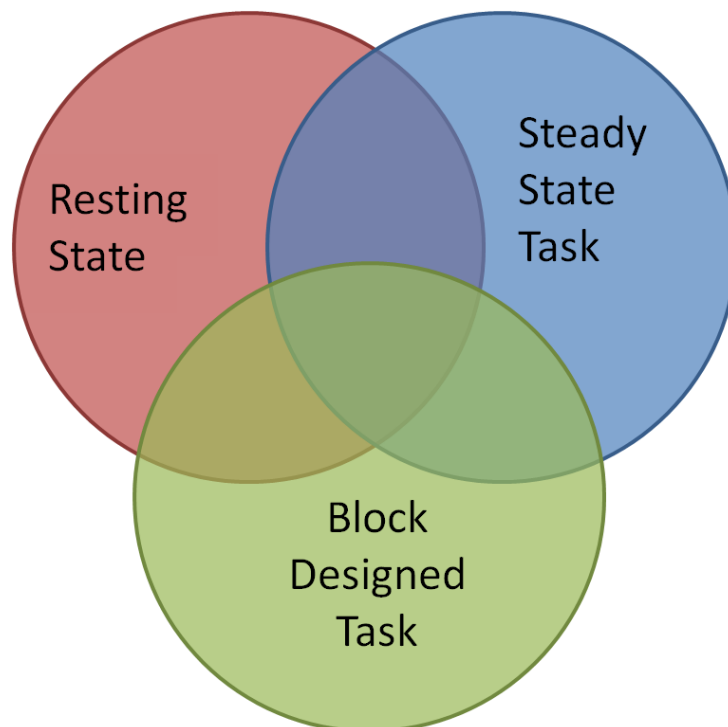


Figure 3.1 - Conceptual architecture of the experiment. The red and blue circles refer to sets of brain regions identified as showing BOLD → ANS relations during the resting- and steady- state epochs. The green circle refers to a block – designed version of the task. This dataset would be analyzed through a comparison of the BOLD activity with the task model.

1. fast-TR resting state epoch: A fast-TR resting state epoch was included to identify regions showing BOLD->ANS relations during rest and to be able to cross-reference these against regions found in the task context. We expected that such a brain system should to some extent resemble the neurovisceral integration one, described by Thayer and colleagues (Thayer et al., 2012; Thayer & Lane, 2009).

2. fast-TR steady state cognitive task: We presented a steady-state cognitive task, mainly for two purposes: one was to test whether there was a common set of areas showing relevant BOLD – autonomic activity effects both at rest and when the brain is involved in performing stressful activity. Areas that will show BOLD → ANS relations in both the resting-state and cognitive task epochs would form part of a very general system related to neurovisceral exchanges. On the other hand, cognitively stressful tasks (like, for example, mental calculation ones) are known to elicit autonomic arousal (Critchley et al., 2003) and consequently to increase the heart rate variability. In this way, the task epoch could evoke more ANS signal variation, with potentially stronger correlation with the BOLD signal.

3. Slow-TR block design: The block design included the same cognitive task as the fast-TR steady state task. It quantified whole-brain activation and deactivation. It also allowed us to quantify the relationship between BOLD and ANS activity during our task when quantified at a conventional TR.

There is a number of ways to characterize autonomic nervous system fluctuations: for example skin conductance response, a slowly-varying indicant related to the level of skin moisture, influenced by the arousal level (Fan et al., 2012; Nagai, Critchley, Featherstone, Trimble, & Dolan, 2004), and cardiac – related indicant, like the heart rate variability (HRV)

(G G Berntson et al., 1997). For the current investigation, we chose to analyze heart rate variability because it is simple to quantify from data acquired within the MR experimental setup, it is non-invasive and it carries substantial information that is known to relate to behavioral effects: Thayer and colleagues (Thayer & Sternberg, 2006) point out that a variability in the length of the beat-beat intervals is a result of a dynamic balance between sympathetic and parasympathetic controls on the heart. They state that heart rate variability, beyond being a mere index of healthy heart function could provide an index of how the brain is giving a dynamical control over the periphery. Heart rate variability denomination refers to several quantities related to ongoing variations in the time of occurrence of heart beats. The basic unity considered is the beat to beat interval, and over this several other quantities are derived. These quantities are mainly classifiable in time domain and frequency domain ones. Frequency domain methods are those involving the study of the spectrum of beat – to – beat interval time-series. They contain information about autonomic activity in that the power of distinct frequency bands (mainly: so-called high and low frequency bands) is related to specific autonomic effects (Malik, 1996). Respectively, high (around 0.25 Hz) frequency power is related to vagal – parasympathetic activity, low (around 0.1 Hz) is somewhat related to parasympathetic, slower sympathetic and even non-autonomic processes. Consequently, one of the most frequently used method to characterize changes in autonomic activity using heart – rate variability is to quantify changes in different frequency bands as a function of different experimental conditions (Montano et al., 1994). Change in the proportion of energy in different frequency bands is taken to refer to the engagement of sympathetic system in order to modulate autonomic activity. Similarly, a heart rate variability sample filtered in different frequency bands carries specific autonomic information: suppressing heart

variability components in a specific band, for example by smoothing, would result in a sample where the contribution of autonomic information relative to that band is lowered. Smoothing using relatively long time windows (e.g.: 8 seconds) would filter out faster contribution, including those related to vagal modulation of autonomic activity, characterizing the sample as sympathetic - related. Using shorter time windows (e.g. 4 or 2 seconds) would on the other hand leave intact the vagal contribution, and return a sample characterized by both the pieces of autonomic information. This holds for time – domain heart rate variability characterizations as well. Root mean square of successive differences (RMSSD) is one of those, related to how the mean heart speed changes in time. This specific time-domain measure is mainly correlated with vagal activity, but can also reflect sympathetic influences (Gary G Berntson, Lozano, & Chen, 2005). Low – pass filtering RMSSD-like sequences using different time – windows, as previously described, should differentiate autonomic system effects they reflect: this is why we decide to choose a quantity related to this time – domain method and manipulate it using moving window averages as our way to quantify heart rate variability.

3.2 Methods

3.2.1 Considerations underlying determination of imaging protocol

The main part of the study consisted of a functional EPI acquisition protocol with a fast repetition time (TR = 400ms). Due to the limited number of slices allowed that can be acquired in this time period, both the task and rest conditions were repeated twice: one with a field of view covering more inferior axial slices (in Talairach space: Z = ~ -20 to ~ 10 mm), and one covering more superior slices (Z = ~ 10 to ~ 35 mm). These two fields of view were prescribed by V.I for each participant at the beginning of the scan, and figure 2A shows the

overlap of these fields of view, for all participants, in common space (see below for information on FreeSurfer based alignment to common space).

A fast TR protocol was used to identify directional relations between BOLD and ANS fluctuations, i.e., to identify temporal offsets between fluctuations in BOLD and autonomic activity. Cardiac events occur at a rate of 0.8-1.2 Hz on average; a typical BOLD sampling frequency used in fMRI studies (TR= 2-3 sec) would not be able to identify small offsets between BOLD and ANS fluctuations and will represent those as simultaneous due to the slower sampling rate. A simulation demonstrating this fact is presented in section 3.3.3.

Second, the range of natural frequencies spanned by cardiac events is unlikely to be less than 0.8 Hz (Malik, 1996), leading to a Nyquist frequency of about 2 Hz. A conventional FMRI scanning protocol has a sampling frequency of about 0.5 Hz, consequently suffering of aliasing of the effects of physiological fluctuations within the time-series. In the literature this phenomenon is well-known and is typically considered from the perspective of treating correlations with ANS as physiological noise. That is, given the lack of ability to identify a directional relation (as explained above), BOLD time series where aliasing occurs are candidates for retrospective correction procedures that remove the impact of ANS activity on the BOLD time series. Retrospective (Glover et al., 2000) methods address the BOLD:ANS correlation problem via Fourier decompositions of BOLD signal. Physiological cycles are considered as quasi-periodic and differences between the image acquisition times and the nearest physiological events are interpreted as phases. In this straightforward way FMRI data could be sufficiently cleaned from noise induced by physiological movements.

However, precisely describing temporal interrelations between BOLD and autonomic activity calls for different solutions. De Luca and colleagues (De Luca, Beckmann, De Stefano, Matthews, & Smith, 2006), for example, explored capabilities of unsupervised methods to characterize resting state networks, and drew on an extremely fast TR (120ms) to clearly demonstrate robustness of network detection, avoiding the aliasing artifact. However, as pointed by Birn (Birn, 2012), this type of solution necessitates a very limited spatial field of view extent: only a few slices could be acquired in order to obtain a satisfying amount of signal in such a small time. De Luca and colleagues (De Luca et al., 2006) in the fast-TR setup, acquired only a single slice.

In the current study we dealt with the Nyquist constraint by setting the TR of our EPI scans at 0.4 s. Pilot tests on the local scanner (4-Tesla, Bruker Medical, Ettlingen, Germany), using a phantom indicated a very good temporal signal to noise ratio (tSNR, evaluated as the ratio between temporal mean and temporal standard deviation, $tSNR = 200$) in an acquisition protocol with 7 slices of 2.1mm. Preliminary tests also suggested 20 seconds of dummy-scan time to reach steady-state magnetization and minimize T1 effects on the functional data.

3.2.2 fMRI imaging parameters

3.2.2.1 Structural scan parameters:

One structural scan was obtained per participant at the beginning of the experimental session, using a 3D T1-weighted Magnetization Prepared RAPid Gradient Echo (MPRAGE) sequence (TR/TE = 2700/4.18 ms, flip angle = 7° , voxel size = 1 mm isotropic, matrix = 256 \times 224, 176 sagittal slices)

3.2.2.2 Fast-TR parameters:

The imaging parameters for Echo Planar Imaging (EPI) functional acquisitions were: TR/TE=400/20 ms, flip angle=39°, voxel size = (3x3x3 mm, 0.45 mm spacing), matrix = 64x64, 7 axial slices with an orientation based on the AC – PC convention and adjusted to follow specific anatomical features of the participants, in order to cover Insula, most rostral portion of the anterior cingulate cortex (ACC), temporal lobe, and posterior cingulate cortex (PCC). We acquired 850 scans (340s) for 2 (one for each FOV) resting-state sessions and 4 (2 for each FOV) steady – state mental calculation task (described below) sessions, 450 volumes (180s) each.

3.2.2.3 Slow-TR parameters

In the slow-TR sessions we preserved the slice orientation used in the fast-TR protocol and acquired 115 whole-brain EPI volumes (253s), using a conventional-TR fMRI setup (TR/TE = 2200 / 33 ms, flip angle = 75°, voxel size = 3x3x3 mm, 0.45 mm slice spacing, matrix = 64x64, 37 axial slices) for each session. One block was a resting state scan (253 sec) and the other a block-design version of the mental calculation task presented before with 4 task/rest cycles in the run (each cycle consisting of 30 sec task, 12.2 sec rest)

3.2.3 Participants

Thirteen right-handed participants (9 males, mean age = 24, SD = 3 years) took part in the study after undergoing a medical screening and signing an informed consent. None of them reported history of illness related to the purposes of the study.

3.2.4 Procedure

Outside the scanner, before starting the session, we instructed the participants on the structure of the session to facilitate task execution during the experimental session. The timeline of the experiment is presented in figure 3.2.

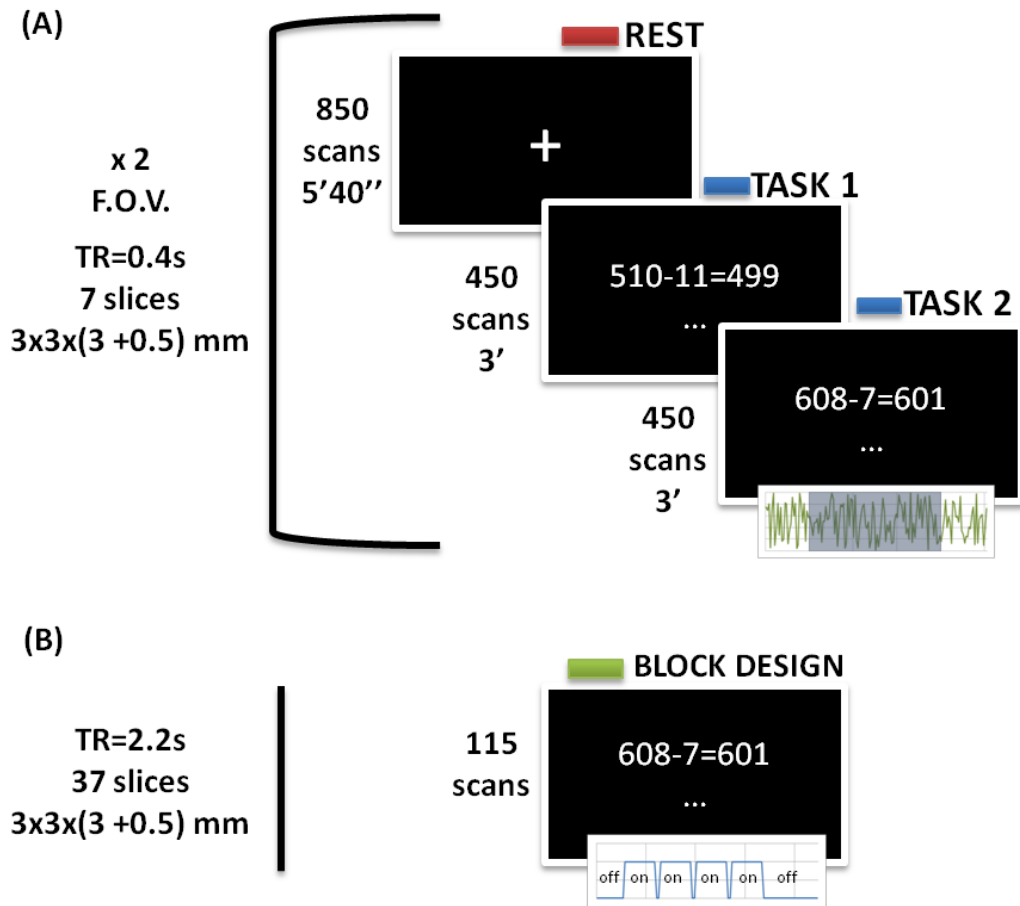


Figure 3.2 - Timeline of the whole experimental session. (A) In the fast – TR sessions we presented the participants with two resting state sessions (one for each field of view), where they were instructed to passively fixate a crosshair put in the center of the screen. We presented four steady – state task sessions (two for each field of view), where the participants were instructed, after a start instruction, to mentally subtract the same number from an initial one and stop when it was notified, 100 seconds later. (B) Block design session consisted in 4 short (30 seconds) blocks of the same task as before.

The orientation of the limited FOV slices was set after acquisition of the structural scan and was propagated to all fast-TR scans in the session

During the task-based 2 fast-TR epochs (180s each; subscripts 1,2 in figure 3.2A) we asked participants to perform a continuous mathematical task covertly without any overt response. After 20 seconds of dummy-scan acquisition times and an additional 20 seconds of rest (passive fixation of a white cross-hair over a black background) a cue appeared to begin the mathematical task. This cue was an arithmetic expression such as “ $510 - 11 = 499$ ” indicating to begin subtracting "11" from the starting number "510". After 2 seconds the cue disappeared and participants were told to continue subtracting covertly until a STOP instruction appeared on the screen (100sec of task), at which point they rested, in the same way as before, for another 40 seconds. To avoid excessive practice effects there were 4 different versions of the subtraction task: continuous subtraction of 11, 7, or 13, and another block where participants subtracted 2 and 3 repetitively (e.g., $510 - 2 - 3 - 2 - 3 \dots$). The order of these 4 variants in the study was assigned randomly to each participant.

During the conventional-TR session (253s), we asked participant to perform the same mental subtraction task. However, in the conventional-TR session they performed the experiment in a block-designed fashion, starting with 10 seconds of rest and then undergoing 4 blocks of subtractions (seeing instructions similar to those presented in the steady-state sessions) each lasting for 30 seconds. The inter-block interval was 10sec and the last block was followed by 40sec of rest.

Participants were scanned during 3 resting state epochs (2 during the fast- and 1 during the conventional- TR sessions), where they were asked to passively fixate a white cross-hair over a black background.

3.2.5 Physiological acquisitions

Cardiac and respiration data were acquired using the scanner's built-in equipment. Cardiac sequences were registered through a photoplethysmograph placed on participants' left forefinger. Respiration data were collected using the displacement of a sensor placed on a belt around participants' chests. Data were sampled at 50 Hz and stored into text files.

Berntson and colleagues (G G Berntson et al., 1997) classify time-delays for autonomic responses according to the nature of the heart rate modulation, defining a range from 1 sec (vagal / parasympathetic activity) to 4 sec (sympathetic activity). A study using vagal or sympathetic nerve stimulation in animals (Kawada et al., 1996) reports differences in the time to reach the maximum response to an induced stimulation, according to different nerves. Vagal stimulation shows maximum effect at about 3 seconds, considerably faster than the effect induced by sympathetic one (about 8 seconds). Spectral properties are also reviewed by Elghozi and Julien (Elghozi & Julien, 2007): low-frequency (0.1 Hz, 7 to 15 seconds) oscillations of heart rate contain information referring to both sympathetic and parasympathetic activity. Faster (about 0.25 Hz, 2 to 5 seconds) oscillations on the other hand contain information related to respiratory sinus arrhythmia (RSA), mediated by parasympathetic activity. Typically, autonomic information from cardiac sequences is derived from relatively long time series, often greater than 300 sec (subsuming on average 360 cardiac events). However, recent work shows reliable assessment over shorter periods. Nussinovitch and colleagues (Nussinovitch et al., 2011) analyzed time-domain autonomic

indicants extracted from time-windows of different lengths. They acquired 300s of cardiac sequences and then compared two parameters extracted from those over time-windows of different lengths (10, 60 s). Parameters were the standard deviation of N-N intervals (SDNN), that practically is the standard deviation of a sequence given by the inter – beat time intervals and the root mean square of successive differences (RMSSD), that is the root mean square of the same sample. They concluded that RMSSD is reliable even when extracted from approximately 12 cardiac events (10 s). Similar results had been obtained by McNames and colleagues (McNames, Thong, & Goldstein, 2003) and Schroeder et al. (Schroeder et al., 2004), the latter suggesting averaging across sessions as a way to improve the significance of the final information. In our current analyses, we take 2 seconds windows of cardiac activity as the basic unit of analysis.

3.2.5.1 Extraction of autonomic-related information:

Initial extraction of the autonomic-related information from the physiological data was conducted as follows: cardiac time-series were first corrected for artifacts (e.g. systematic spikes due to the acquisition device amplitude range). Local maxima (peaks) in the cardiac time series were identified. Candidate beats were initially identified statistically as peaks with a magnitude greater than 3 standard deviations from the mean amplitude.




Figure 3.3 - Peaks in the cardiac sequence are automatically identified, but sometimes there are misses that require manual intervention.

This procedure returned a raw time-series containing automatically determined timing of cardiac events, which was then manually evaluated and fine-tuned.

The manual analysis corrected for misses and false positive in the data that could not be identified by the automatic procedure, as shown in figure 3.3.

We then evaluated the inter-beat interval (IBI) calculating the difference between the time of occurrence of two adjacent cardiac events (defined as the peaks seen in the recording):

$$IBI_k = t(B_{k+1}) - t(B_k) \quad (1)$$

We finally obtained a set of couples:

$$\{t_i, IBI_i\}$$

where IBI_i is the inter-beat time interval, and t_i is the timing of the second beat.

This set series presented the well-known problem (Guimarães & Santos, 1998) of varying sampling frequency, i.e. the fact that samples are not equidistant. To perform a time (or frequency) – domain analysis would not be possible using this sequences, as far as the same amount of samples could refer to different temporal windows in different part of the set: for example first 10 samples could fall within 8 seconds, and last 10 within 12. However, given the fact that we know the time of occurrence for all the samples, we could interpolate and obtain a new time-series. We chose as the final sampling frequency the same as the fast-TR one (2.5 Hz) and we filled in the gaps using a spline function. We did this to ensure smoothness of the final sequence and to avoid problems coming from non-piecewise defined functions like linear or plainly polynomial ones. We ended up with one IBI time series aligned and ready to be compared with FMRI data.

3.2.5.2 Behavioral effects on the heart speed:

By definition, inter beat interval vectors amplitude decreases when heart start beating faster, and increases when it slows down. To study cognitive effects on the heart speed we considered IBI within a period of 72 seconds taken within the central part of the resting-state epoch and during the central part of the steady-state task execution period. We reported mean beats-per-minute and the differences between the conditions. We also evaluated the standard deviation of the IBI and compared changes in the two epochs.

3.2.5.3 Time-lagged correlations with proxy vectors:

Point-by-point correlations account for simultaneous effects occurring in the considered samples. These correlations identify structural similarities between two vectors, for example time-series. Granger causality and other approaches for effective connectivity or lagged correlation account for time delayed factors when quantifying synchronization. Our goal in

the current analysis, however, is to identify and quantitatively describe potential relationships between fluctuations in one time-series (BOLD) and the variance of a set of values in the other time series, all of which occur in the near future. To this end we converted the raw measurements to proxy vectors summarizing blocks of cardiac activity in windows of different length. A statistically significant link between the BOLD time series and this derived sample would be an index of non-simultaneous effect between the two elements. For instance, the regression model where the mean heart rate in samples $t+1.. t+10$ is modeled by BOLD activity is timed T has the form:

$$BOLD = \beta * MEAN(IBE_{t+1,t+2}, IBE_{t+2,t+3}, \dots, IBE_{t+10,t+11}) + \varepsilon \quad (2)$$

which means that a voxel time-series (left hand of the equation) is explained by a RMSSD-like sequence consisting of the mean of the inter-beat intervals in the next 10 samples (right sign of equation). A similar analysis explaining bold by the *variance* of IBI would have the following form:

$$BOLD = \beta * VAR(IBE_{t+1,t+2}, IBE_{t+2,t+3}, \dots, IBE_{t+10,t+11}) + \varepsilon \quad (3)$$

The beta coefficient that is returned by the regression, i.e., its sign (positive or negative), its magnitude and its statistical significance are all voxelwise quantities composing directed correlation statistical maps.

A RMSSD – like sequence composed by differences in time of occurrence of heart beats relates mainly to vagal activity, but also carries sympathetic effects (Gary G Berntson et al., 2005). To cover a wide range of potential generators of fluctuations we chose four different time-windows in which the heart-rate variability was quantified, effectively implementing 4 different low-pass filters with different cut-off frequencies. (.5, .25, .125, .025 Hz). The

length of these sliding windows was 5, 10, 20 or 40 volumes (at 0.4 sec /volume) corresponding to 2, 4, 8 or 16 seconds or cardiac variance. This sliding window averaging results in cutting oscillations in frequencies higher than respectively .5, .25, .125, .025 Hz. These frequency bands reflect a subdivision of heart rate spectrum, where higher bands refer to vagal activity and the lower to the sympathetic one. Maps from different proxy-vectors would refer to different generators of potential relationships between BOLD and autonomic activity fluctuations.

3.2.5.4 Spectral characterization of smoothed sequences

To show how different proxy vectors carry different spectral information, we performed an analysis of the amount of power in the heart rate frequency bands, in each of the sliding window lengths we had used to smooth the heart rate variability sequences.

We analyzed exactly the same time – series we inserted in the BOLD – AA regression: heart rate variability maps interpolated at 2.5 Hz, z-scored and finally smoothed over 5, 10, 20 and 40 samples (referring to, respectively, 2, 5, 8, 16 seconds). Here, in addition to this, we also considered non-smoothed plain raw heart rate variability sequences (evaluated using eq 1), which serves as a pre-filtering reference. We considered the central 750 (out of 850) samples of the time-series and we evaluated their power spectrum using squared absolute value of the Fast Fourier Transform. We then normalized the magnitude of the spectrum dividing by its maximum value, to obtain an index of the fraction of power belonging to each frequency.

We then defined four rectangular frequency bands, 0.05 Hz wide, centered around 4 frequency values referring to the sliding – window lengths: the highest, ranging from .48 to

.52 Hz, centered in .50 Hz and referring to 2 seconds (5 scans) sliding window; then another from .23 to .27, centered in .25 Hz and referring to 4 seconds (10 scans) sliding window; the third was centered in about .12 Hz, ranged from about .10 to about .14 (8 seconds – 20 scans) and the last ranging from about .04 Hz to about .08, centered in about .8 Hz, referring to 16 seconds (40 scans) sliding window length. These windows refer respectively to very high, high, low, very low frequency bands in heart rate variability spectrum and could be related, especially the two central one, to respectively vagal (up to about .25 Hz) and sympathetic (up to about .12 Hz) modulation of the autonomic activity.

We averaged the power spectrum within these ranges and evaluated the fraction of power they carry, for each person, for each sliding window length and then we averaged the fractions across the participants.

3.2.6 FMRI Data processing

Functional acquisitions for all the experimental sessions (fast- and slow- repetition time ones) underwent the following spatial – temporal preprocessing pipeline: first, from the fast-TR time-series, the first 50 volumes (20 seconds) were discarded, based on preliminary tests quantifying the delay towards steady-state magnetization. Similarly, the first 14 scans of the slow-TR protocol (30.8 seconds) were discarded from the slow-TR time-series. Then, despiking was implemented using AFNI's 3dDespike function so that time-points deviating significantly in value from the median were replaced using an interpolation procedure. Despiked time-series were corrected for slice-timing differences and spatially aligned to a reference acquisition. The alignment procedure (implemented via AFNI's 3dVolreg) also returned 6 vectors representing the trends of six motion parameters. A spatial smoothing

using a tridimensional Gaussian 6-mm isotropic kernel was finally applied to the time-series in order to improve their signal-to-noise ratio.

Statistical maps resulting from all the regression procedures (explained in the following sections) were projected into 2D-surfaces using Freesurfer / Suma processing pipeline as follows: single participant's structural images were aligned to the first functional run (to maximize T1 effect) and alignment was manually checked and adjusted when needed. The next steps were implemented using FreeSurfer procedures. Here, the anatomical-in-the-functional-space is inserted into a process where it is registered to a reference space, segmented and skull stripped. In this way all the non-grey matter could be carefully isolated and removed. The final result is a well-formed gray matter surface ready to be inflated.

3.2.6.1 Slow-TR analysis and ROI extraction:

Each participant's slow-TR time-series were scaled by their temporal mean and multiplied by 100 in order to obtain mean-normalized time series that could be combined across participants and allow group-level analyses of BOLD signal change data. Statistical parametric maps of task-induced effects on the brain function were produced through a voxel-wise regression of the scaled time-series against a task-model regressor.




Figure 3.4 – Timeline of slow-TR, blocked-design task. Blue and Red curves refers to the models of, respectively, start and stop instructions. Green ones model the actual task period. There were 4 on/off cycles of continuous mental subtraction. Volume acquisition time (TR) was 2.2 sec and we ended up with having about 220 seconds of relevant acquisitions.

The regressors and design of a given participant are given in figure 3.4. As shown in the figure, there was one regressor capturing the visual instruction to 'Start' and 'Stop' the calculation task, and one regressor capturing the task period itself.

We convolved these boxcar regressors with a canonical hemodynamic response function. We then inserted these three regressors in a design matrix including also the trends of 6 motion parameters, considered nuisance variables. For each participant, this regression produced statistical parameter maps identifying regions with above- and below-baseline activity during the subtraction task. The resulting maps were then then projected from 3D-volume domain to 2D-surface one, using Freesurfer / SUMA processing pipelines.

These beta values on the 2D surface were inserted into a whole-brain voxel-wise second-level group analysis. This test evaluated, for each surface vertex, whether the mean Beta value departed from zero. The correct threshold was calculated using cluster-based family-wise error (FWE) following the Monte-Carlo simulation method suggested by Forman and

colleagues (Forman et al., 1995). Thresholding procedure returned surviving areas with positive or negative relationship with respect to the “ongoing” condition. These effects are commonly referred to as “deactivations” and “activations”. In this analysis, the single voxel significance level was $p < .01$ (uncorrected), controlled for family-wise error using cluster-wise correction ($p < .05$). The clusters returned by this procedure further served as functional ROIs for further analyses.

3.2.6.2 tSNR of fast- and slow-TR dataset:

Prior to analyzing the fast-TR data we compared the temporal signal noise ratio within that protocol to that found in the slow-TR one. tSNR was quantified as the inverse of the coefficient of variation: the temporal mean divided by the temporal standard deviation. This value was calculated on the raw data prior to any processing (no slice timing correction, motion correction, alignment or smoothing) and for this reason should not be taken as indicative of tSNR in the data processed downstream, but as an indicator of initial data quality. This quantity was evaluated voxelwise for both fast- and slow- TR, considering the same amount of volumes (N=100), and then spatially averaged for each dataset.

3.2.6.3 Fast-TR lagged correlation analysis:

First, Fast-TR steady-state task and resting state datasets underwent a scaling identical to the one performed on slow-TR ones. To identify lagged correlations between BOLD and cardiac HRV *in the resting-state data*, the resting state time series were regressed against one autonomic activity proxy vector. Out of the 800 volumes, first and last 50 scans were discarded to avoid artifacts caused by initial and final transients in the proxy vectors.

To identify lagged correlations between BOLD and cardiac HRV *during the active task performance*, the task-data collected during the fast-TR protocol were processed as follows.

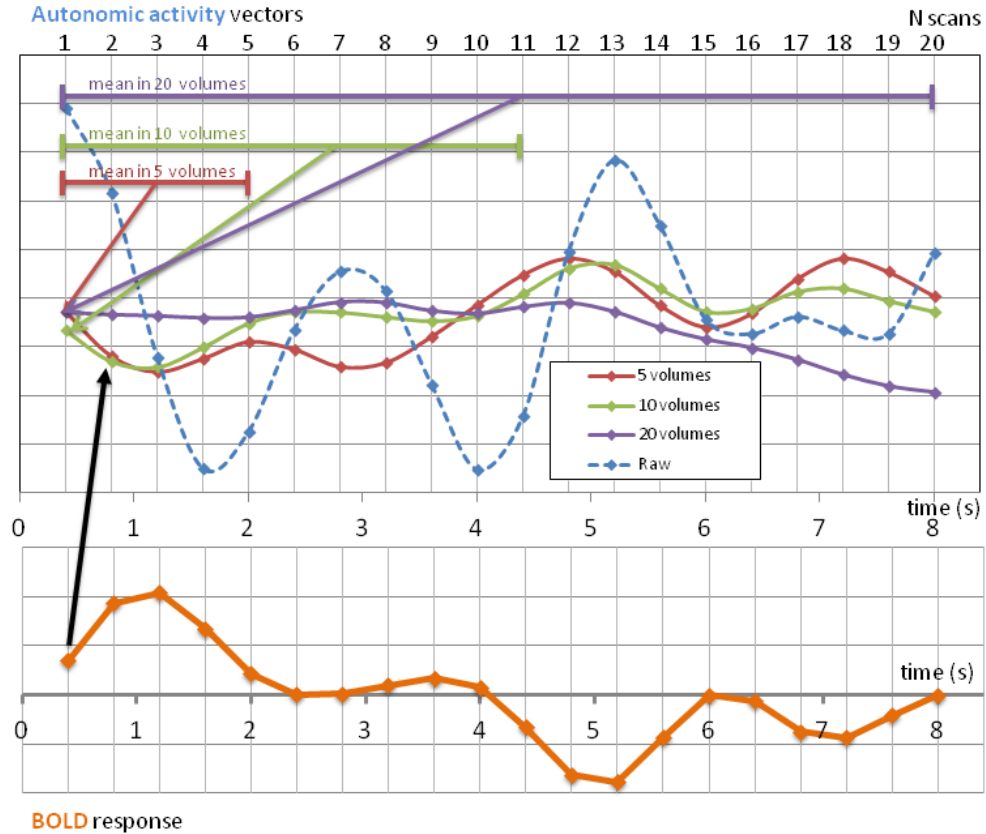


Figure 3.5 - Sliding – window lagged correlation method. On top there is the original (blue dashed line) inter – beat interval sequences, referred to as “raw” in the text. We report both the time- (bottom) and the fMRI fast sampling rate (top) on the x axis. Red, green and blue curve refers to the sliding – window averages of the blue one. This is the way in which we built our autonomic activity proxy vectors. At the bottom we show (in orange) BOLD response as well, and the black arrow refers to the regression procedure we introduce to obtain maps of directional interactions.

These steady-state task time-series were 400 volumes long overall. However, the initial time period within each epoch consisted of instructions, so that participants performed the task for only for 250 (100 seconds). Of these 250 functional volumes, we discarded first and last 35 volumes (28 seconds overall), ending up with 180 (72 seconds) volumes. These time series of length 180 were regressed against the ANS proxy-vector, as we show in figure 3.5.

Statistical maps obtained from this regression were then projected to the surface using AFNI's 3dVol2Surf utility.

3.2.6.4 Masks and group analysis of the fast-TR protocols:

While we attempted to set the partial fields of view (in the fast-TR protocol) so that they were highly similar across participants, there were of course differences in areas covered between participants. To establish group maps for this protocol it was therefore first necessary to identify the brain areas covered by a sufficiently large number of participants, and to conduct the group analysis only in those regions. We identified areas that fell in the same field of view across participants in the following way. First, we constructed in the volume domain, intracranial binary masks using a method based on voxel intensity in the functional scan, which indicated the volume space where data was acquired. This was done for each field of view (High – Low), for each condition (Rest – Task1-2).

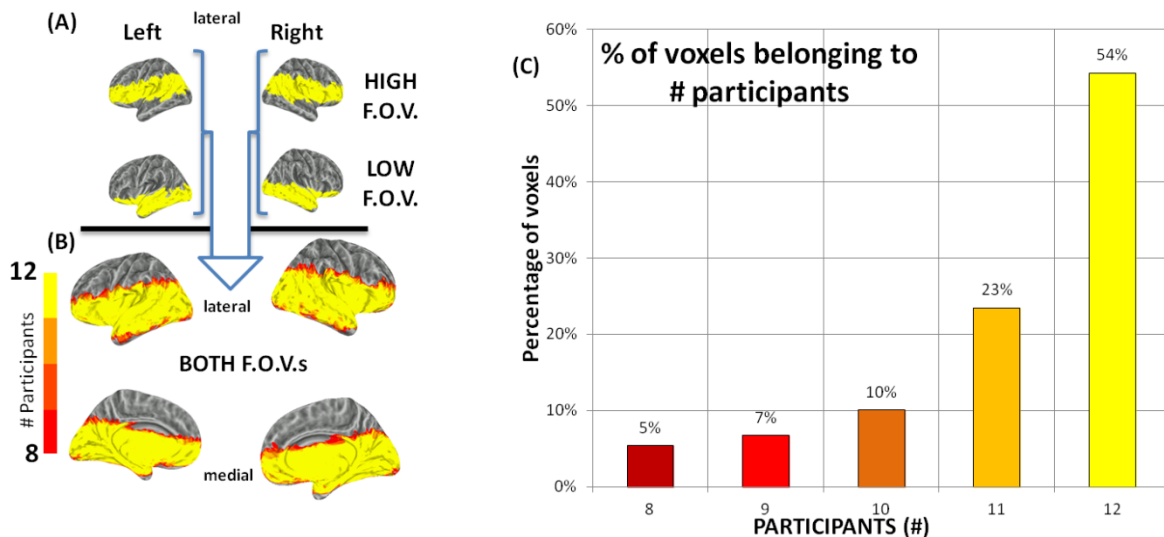


Figure 3.6 – (A) Partial fields of view originally acquired were averaged, merged and combined through the participants in a conjunction map showing which were the most common areas. We produced final masks saving vertices belonging to 8 or more participants. (B) In the final masks most of the voxels (more than 70%) belonged to 11 or 12 participants.

We then projected volumetric masks into 2D-surfaces; in this domain we averaged within conditions ending up with having, for each participant, one resting state and one steady-state task mask for each field of view. For each participant we then merged the masks of the High and Low field scans to obtain two single-participant final masks, one for resting state data and the other for steady state task.

We then created a conjunction mask to identify which regions were sufficiently covered for multiple participants, and selected only those surface vertices where 8 participants at least provided data (see figure 3.6b). As shown in figure 3.6 these covered the insula, the ventral part of the medial prefrontal cortex, as well as occipital and temporal regions. Furthermore, while a threshold of 8 participants was a lower bound, in more than 70% of them are we obtained data from 11 or 12 participants (see figure 3.6c).

We restricted second-level statistics on fast-TR data within these group masks: a vertex-wise single participant's statistical parametric map was performed and the results thresholded using a simulation similar to that described for slow-TR data, but restricted within the relevant field of view.

3.2.6.5 Cross-reference with the task related effects

To investigate the functional characteristics of the areas showing relevant BOLD – AA relationships, we first considered the set of distinct regions obtained by the clustering procedure on BOLD – AA maps. Then we considered each of these areas as a functional region of interest (fROI). In each of these fROI we could then evaluate their activity profile in the slow-TR task to identify if they would be considered as active or inactive in a typical block design employing the same calculation task. In this way we ended up with having

couples of functionally defined indicants, referring to different aspects of the functionality of one area: one related to the relationship to the autonomic activity during an ongoing task, the other to the effect elicited by the execution a block – designed version of the same task.

3.3 Results

3.3.1 Behavioral effects on heart speed:

The mean inter-beat interval was taken as a parameter of interest in this analysis. The parameter was calculated from a 72 sec period in the central part of the fast-TR resting state session and the fast-TR steady-state task. During rest we obtained, for the mean heart beat interval, a value of 70.78 BPMs and 2.24 SEM. During execution of the subtraction task the heart rate increased to 73.68 (2.52 SEM) BPMs, which constituted a statistically significant increase in heart rate [$t(11)=2.51$, $p=.028$]. The standard deviation of the IBI is another measure typically used to assess stress impact (Malik, 1996). This measure showed significant [$t(11)=6.56$, $p < 5e-5$] reduction for task (SD = 4.88 BPMs) vs. rest (SD = 0.70 BPMs) of about 6 times.

3.3.2 Task – induced effects (Slow – TR session)

A whole brain analysis was conducted to identify areas showing above- and below-baseline BOLD signal levels during performance of the calculation task in the blocked Slow-TR session.




Figure 3.7 – Task induced activation and deactivation maps from slow – TR sessions. Statistical coefficients between fMRI data and the task model described by the green curve in figure 3.4 are mediated between participants (N = 12) and showed on the maps, weighted using T – statistics and the clustering procedure described in section 3.2.6.1. Hot and cold colors refer respectively to task – induced deactivations and activations.

We refer to above-baseline active as ‘activation’ and below-baseline activity as ‘deactivation’. Figure 3.7 shows the statistical parametric map of task induced activation and deactivation. In the left hemisphere, execution of the task individuated twelve macro – areas: the biggest one referred to a deactive cluster located in the frontal part of the brain including orbito – frontal areas, anterior cingulate cortex, anterior insula, para-hippocampal gyrus, anterior temporal areas. Left – Posterior cingulate cortex / Precuneus, superior temporal sulcus / angular gyrus, central sulcus and occipital gyrus also showed significant deactivations. On the same hemisphere, we report significant activations in the superior frontal gyrus, superior frontal sulcus, precentral sulcus, inferior frontal sulcus, circular sulcus of the insula, post – central, intraparietal and and calcarine sulci. Right hemisphere clustering

first returned an extended deactivated area covering temporal areas, parahippocampal gyrus and anterior insula. We also obtained deactivations in Anterior cingulate cortex / orbitofrontal areas, posterior cingulate cortex / Precuneus, superior temporal sulcus, orbital lateral areas, central sulcus. Right hemisphere activations were located within intraparietal and superior frontal sulci, circular sulcus of the insula, post – central sulcus and calcarine sulcus. These sets of areas are consistent with previous findings for similar tasks (Kawashima et al., 2004), including a laterality effect (left hemisphere shows relevant effects in the precentral sulcus, while right one does not) for activations. Deactivation network show results similar to those found in previous findings (Greicius & Menon, 2004), composing a set of areas resembling that commonly referred to as “task – negative network”. These maps of the effects of slow-TR task are crucial for our purposes: first, they formed a functional signature showing that participants were indeed involved in the execution of the task. Furthermore, as we detail below, these statistical maps allowed us to characterize the nature of the activity of specific fROIs identified as showing bold-ANS relations.

3.3.3 Simulation: apparent synchrony between BOLD and ANS fluctuation due to low sampling rate:

Simulations were conducted to evaluate whether decreasing the temporal resolution of the fMRI time series hinders the possibility of identifying lagged correlations between BOLD and ANS. We simulated (N=1000) couples of vectors (one referring to FMRI and the other to a generic physiological acquisition). The vector-pairs simulated 400 functional volumes (TR = 0.4 sec) and physiological data acquired during the scan. The correlation between them was fixed at Pearson’s correlation ($R = .3$), using a procedure for creating correlated vectors. Finally, the 'ANS' vector was shifted 5 samples (2 sec) forward to simulate a lagged

regression. That is, any correct regression solution would identify that the vectors had a maximal correlation ($R = 0.3$) at a lag of +5.

To represent the effect of acquiring the same intrinsic information using different sampling frequencies we downsampled the BOLD time series by a factor of 2, 4, 5 and 8, thus simulating the time series that would have resulted if the BOLD data were acquired at a lower sampling rate, generating BOLD series that would be acquired with a TR of 0.8, 1.6, 2.0, 3.2 seconds respectively. We then calculated a lagged correlation (cross-correlation function) between the two downsampled vectors; this returned both a correlation and a lag estimate for which the correlation was maximal.

Although the correlation implemented was quite low, a regression model using the original temporal resolution (TR=0.4) resulted in a good fit, and solved for the correct lag in almost all the simulations (98% of simulations estimated $r = 0.3$, 100% of simulations correctly estimated the lag of 2 sec). The effect of downsampling and concomitant loss of temporal resolution lead to a decrease of the estimated correlation. Statistically significant correlations were we obtained in 31%, 19% and 19% of the simulations for the 0.8, 1.6, and 2.0 TR downsampled series. These correlations, furthermore, tended to be maximal at time lags different from the one we imposed (though they were quite accurate at the 0.8 sec TR version). We obtained, as an average of 1000 simulations the following values: 1.8 sec (0.8 sec TR), 1.2 sec (1.5 sec TR) and 1.7 sec (2.0 sec TR). The slowest simulated repetition time (TR = 3.2 seconds) returned an average of maximum correlation of .03 (10% of the initial value), and this maximum correlation occurred on average with an estimate lag of 1.1 sec on average. To summarize, a moderate lagged correlation between two vectors that could be

identified at a lag 5 with a fast 0.4 sec TR could not be identified with high accuracy at lower temporal resolutions.

3.3.4 tSNR of fast- and slow-TR dataset.

We found that tSNR in fast – TR sessions (tSNR = 20.3) was significantly [$t(11) = 3.77$, $p = 0.003$] higher than the slow – TR one (tSNR = 17.5), which showed a value about 16% lower. This suggested that the information contained in fast – TR datasets was valid and suitable for data – analysis.

3.3.5 Spectral characterization of smoothed sequences

As should be the case, different sliding window lengths suppressed contributions from different part of the spectrum and, consequently, the amount of fluctuations related to specific modulations of the autonomic activity.

Starting from the raw version of the heart rate variability, where all the four frequency bands contain a significant part of the power, each sliding window average begins suppressing parts of the power spectrum, leaving only slower processes to characterize 16 – second smoothed time – series.

From the autonomic activity perspective, most relevant sliding window lengths seem to be those smoothing within 4 and 8 sec windows, in that they carry a significant amount of power in the two central bands, which are known to refer to sympathetic and vagal modulation of the autonomic activity. Slower integration windows, in any case, still show relevant fractions of power belonging to relevant (e.g. low – frequency, .12 Hz) bands.

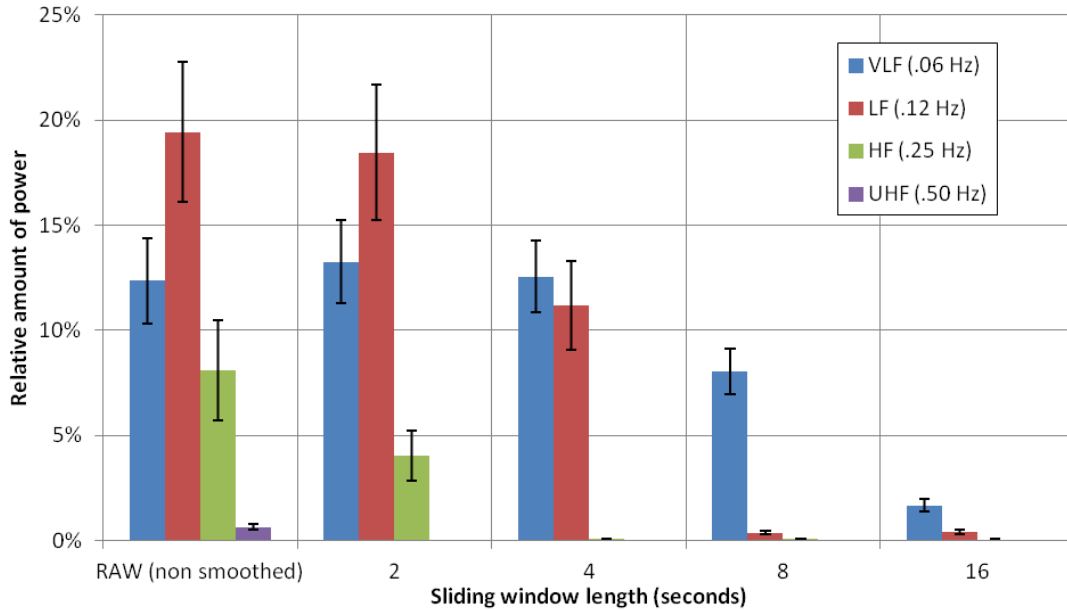


Figure 3.8 – Effect of the smoothing on spectral components of our heart – rate variability time – series. Different sliding windows suppress contributions from different parts of the spectrum.

3.3.6 **BOLD - Autonomic Activity maps:**

The task and rest periods in the fast-TR scans were analyzed to identify regions showing a lagged relation between BOLD activity and subsequent Heart Rate, as described in section 3.2.6.3 of the methods. These analyses could identify regions where increased BOLD predicted higher ANS reactivity (higher hear rate) or predicted lower ANS reactivity (lower heart rate). Therefore, the direction of the correlation is an important part of the result.

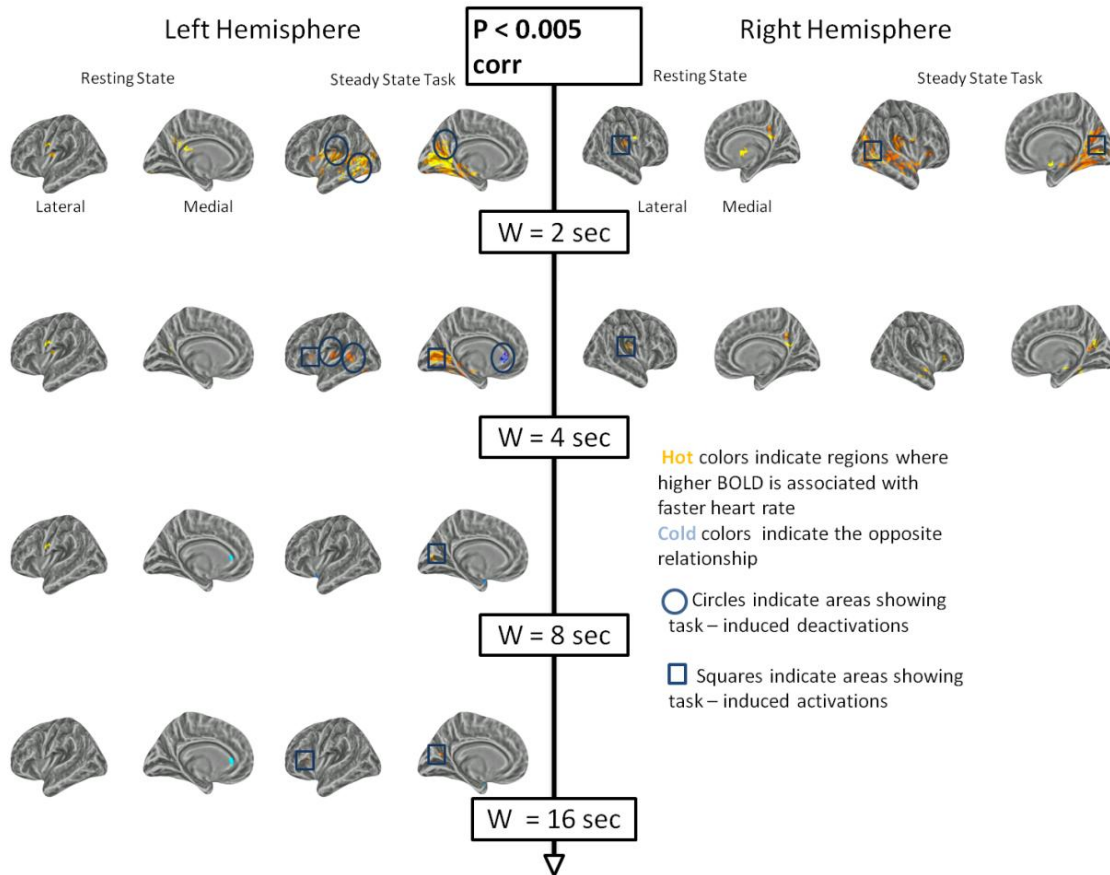


Figure 3.9 – BOLD – autonomic activity maps at different sliding windows, from top (shortest, 2 seconds) to bottom (longer, 16 seconds). Left and right panels refer to respective hemispheres. Within the panels, on the left there are shown resting – state results, on the right those from steady state task. Circles and squares refers to results from cross – referencing procedure explained in section 3.2.6.5.

For the resting state activity, we identified a number of clusters with positive correlations. This predominance is complete when averaging the autonomic activity using short (2 – 4 seconds) sliding windows (i.e., without cutting almost any frequency band). The set of significant areas included bilateral superior / posterior insula, subcentral gyrus, parietal / occipital and posterior cingulate cortices. Increasing the sliding – window length result in removal of some regions.

FAST - TR RESTING STATE																
Location	Extension (mm2)								Mean T Value							
	Left Hemisphere				Right Hemisphere				Left Hemisphere				Right Hemisphere			
	2	4	8	16	2	4	8	16	2	4	8	16	2	4	8	16
Posterior Insula	116	45			74	118			3.4	3.3			3.6	3.3		
Subcentral Gyrus	67	99	82		62				3.5	3.6	3.5		3.6			
Posterior Cingulate Cortex	42				154	59			3.7				3.8	3.7		
Posterior / Parietal	32	38				70			3.4	3.4			3.3			
Anterior Cingulate Cortex			46	72							-4.1	-4.0				
Occipital / Lingual	39								3.4							

Table 3.1 – Table for BOLD – AA effects in the resting – state condition. Left column reports spatial information about where the cluster is located. Central column contains information about the extension of the cluster (in mm2) and the right column describes the statistical power (N = 12 participants) and the sign of the BOLD – AA relationship found in the cluster. Clusters were sorted according to their extensions, considering both the hemispheres.

Implementing a 4 seconds sliding window map identified positive effects in bilateral posterior insula and posterior cingulate cortex. Using longer sliding windows we did not find significant results in the right hemisphere. On the left, we found a positive cluster in the subcentral gyrus (8 seconds) and a negative cluster located nearby the anterior cingulate cortex. This region showing a negative effect is also the only one surviving at very long time – windows (16 seconds). These results could be compared to those obtained by Fan and colleagues (Fan et al., 2012), where correlation maps between fMRI (using a conventional repetition time, TR = 4050ms) and skin conductance at rest were examined. Positive correlation clusters were found there in calcarine sulcus, posterior temporal areas and anterior insula. A negative correlation was found, among other regions, within the anterior cingulate cortex area, resembling what we found with our experimental setup. On the other hand this analysis also revealed negative correlations in other parts of the brain like posterior

cingulate cortex where we noticed positive correlations between BOLD and arousal.

FAST - TR STEADY STATE TASK																												
Location	Extension (mm2)								Functional Behavior (from block design)								Mean T Value											
	Left Hemisphere				Right Hemisphere				Left Hemisphere				Right Hemisphere				Left Hemisphere				Right Hemisphere							
	2	4	8	16	2	4	8	16	2	4	8	16	2	4	8	16	2	4	8	16	2	4	8	16	2	4	8	16
Posterior / Parietal areas	3558	53			1358	291			N/A	N/A			A				4	3.48			3.61	3.6						
Intraparietal	819				670	261			N/A				N/A				3.84				3.92	3.7						
Temporal / Occipital	75				1188	206			N/A				N/A				3.85				3.6	3.68						
Calcarine Sulcus		1157	65.5	70	76					A	A	A	N/A				4.11	4.04	3.55		3.35							
Anterior Insula	883	248			126				D	D			D				3.67	4.07			3.62							
Sup Temporal Areas	158	282			745	67.7			N/A	N/A			D	D			3.34	3.66			3.76	4.12						
Posterior Temporal areas*	1174				34.5				N/A				N/A				3.99								3.68			
Posterior Insula	479	51.2		33	205	86.5			N/A	N/A		A	N/A	A			3.54	3.33		3.31	3.31	3.42						
Lateral Fissure					679								A								3.49							
Inferior Temporal Gyrus		63.5			557					N/A			N/A								3.68				3.7			
Posterior Cingulate Cortex	413								D								4.1											
Occipital / Lingual		409								N/A											3.73							
Planum Polare			297	108							D	D									-4.19	-3.53						
Middle Occipital Gyrus *	301				33				D				N/A				3.72								3.59			
Superior Occipital / Transversal		202								N/A											3.71							
Subcentral Gyrus	104				36				D								3.51								3.47			
Transversal collateral	136								N/A								4.05											
Anterior Cingulate Cortex		78								D											-3.6							
Superior Occipital		50.9								N/A											-3.59							

Table 3.2 - Table for BOLD – AA effects in the steady – state condition. Left column reports spatial information about where the cluster is located. We indicate with a star macro – areas relative to more than one cluster coming from the thresholding procedure. The next column contains information about the extension of the cluster (in mm2). The second last from the left contains information related to the cross – referencing procedure. When the average of the beta – values coming from block – design was significantly positive or negative, we reported A or D, and N/A otherwise. The right column describes the statistical power (N = 12 participants) and the sign of the BOLD – AA relationship found in the cluster. Clusters were sorted according to their extensions, considering both the hemispheres.

Fast – TR steady – state task maps could be described in a similar way, identifying areas showing positive and negative correlations between BOLD and autonomic nervous system. These maps appear to identify more areas than those found in the resting – state maps, and they too show that at longer sliding window-lengths, the right hemisphere did not show any significant effects. Using very short (2 seconds) time – windows, like in the resting – state maps, all the significant effects are ones where there is a positive correlations between BOLD and autonomic activity. There is an extended positive cluster in bilateral occipital / parietal areas; positive effects are also found in the left posterior and right posterior / anterior temporal areas and anterior and posterior subdivision of the insular region.

Using a longer time – window (4 seconds), positive correlations are found in a cluster within the left calcarine sulcus / lingual gyrus area, and in the right occipital / parietal cortex. Left posterior and right anterior temporal areas, left posterior and bilateral anterior insular cortices showed positive BOLD – AA relationship as well. On the other hand, left anterior cingulate cortex, using 4 seconds integration window, showed negative correlations with autonomic activity. At longer sliding – windows only positive effects in left calcarine (8 seconds) and left anterior insula (16 seconds) survived.

3.3.6.1 Cross – reference with task – related effects

The next analysis step aimed to identify whether the regions showing BOLD -> ANS relations were ones typically activated or deactivated by the task. To implement this analysis each of the clusters identified in the BOLD -> ANS analyses was considered as a functional region of interest, and the mean activity level in that cluster in the slow-TR task was calculated.

This analysis would be particularly important if it were to identify that some of these fROIS were considered ‘deactive’ during task performance.

A few regions showed a relatively intuitive result pattern. These were regions showing a positive correlation between BOLD and later autonomic activity (in the fast-TR task), and that were categorized as showing above-baseline activity as estimated by the block design part of the study. These regions included the calcarine sulcus and anterior insula.

However, a number of regions showing a positive correlation between BOLD and ANS activity were actually categorized as *deactive* in the block design task. Areas showing this

pattern included bilateral posterior insula, left posterior cingulate cortex, middle occipital and subcentral gyri.

The region where we found a negative relationship between BOLD and ANS (the anterior cingulate) was found to be deactivated by the cognitive task. That is, on the whole the regions was associated with below-baseline activity, less deactivation (higher signal) was associated with lower downstream ANS activity.

This analysis revealed an interesting functional segregation of the insular region: while both anterior and posterior parts (segregated by our clustering procedure) showed positive BOLD -> ANS correlations, they responded differently to task performance itself, showing respectively activations and deactivations. This could refer to already known insular segregation results (Deen, Pitskel, & Pelphrey, 2011), but also to more sophisticated theories of interoception [i.e. the representation of the physiological condition of the body, (Craig, 2009)] assigning different roles for the two insular subsystems, even if both related to neurovisceral integration problem.

3.4 Discussion

Goal of this work was to address three general questions: first, are there brain areas where cognitive activity is followed by systematic changes in autonomic activity. Second, to the extent that such are found, do they perform similar functions during task and rest states. Third and last, are the brain regions showing this relationship ones that are considered active by task, or do they also include brain areas deactivated by the task.

3.4.1 BOLD - Autonomic Activity regions

One the most important result of the current study is a technical one – that it is possible to use fMRI to identify areas where the relation between BOLD fluctuations and near-future ANS ones is causally linked, and that the use of fast-TR method reveals a lagged relation that would appear as synchronous if a regular TR were used. For the neuroimaging community this implies that regions where BOLD activity appears to correlate with ANS fluctuations are not necessarily ones that are artifactually fluctuating with cardiac or respiratory fluctuations, but actually regions driving those fluctuations.

We identified a consistent set of areas showing significant positive BOLD - autonomic activity effects including occipital -parietal cortices, anterior and posterior temporal areas, anterior and posterior insula, calcarine sulcus and posterior cingulate cortex. We also found that anterior cingulate cortex activity is negatively correlated with autonomic one. Our results are consistent with a recent argument made by Thayer and colleagues (Thayer et al., 2012) on the nature of autonomic regulation. In that work, they propose an explanation of top – down modulation of autonomic activity, suggesting a detailed hierarchical model of regulation, shown in figure 3.10. As ‘top’ modulators of ANS activity they include anterior and posterior parts of the cingulate cortex and anterior insular cortices. They furthermore performed a meta – analysis of 8 works (N participants overall = 191) whose goal was to localize brain structures related to heart rate variability. The meta – analysis was performed to see whether a fine – tuning of the functional segregation of ACC / ventromedial prefrontal cortex areas was possible. They showed an interesting subdivision of areas within the region, where peri – gential anterior cingulate cortex turned out to be a subdivision present in all the studies. This specific part of the ACC is exactly the one where we obtained our results.

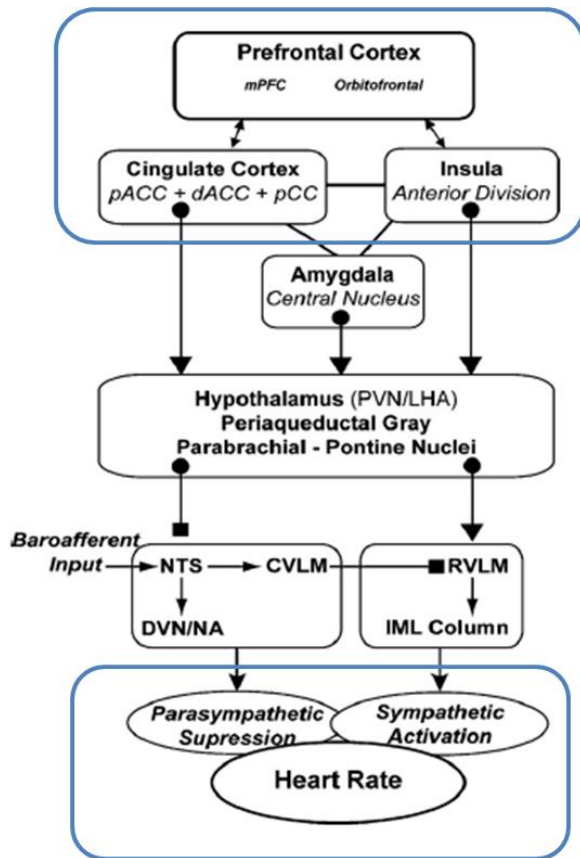


Figure 3.10 – [Taken from (Thayer & Lane, 2009)] – A schematic representation of top – down modulation of heart rate from the cortex to sympathetic and parasympathetic effects. Blue boxes emphasize those subsystems we were able to analyze with our data: the cortical ones on top, and the sympathetic and parasympathetic effects related to heart rate variability and its spectral components.

We reported in our results clusters of significant effects in anterior and posterior cingulate cortex and anterior division of insula. This could suggest that we were able to reconstruct the network of areas involved in this cortical module of neurovisceral regulation.

The same part of ACC was also described in a work by Wager and colleagues (Wager, Waugh, et al., 2009) where they combined FMRI and heart rate acquisitions, presenting the participants with a social threat task. They show that activity in the peri - gential anterior cingulate cortex was negatively correlated with heart speed. The contribution of the current work to these prior models and findings is indicating the directional relation between ACC activity and heart rate features.

Our findings indicate that in general, BOLD ->ANS relations are stronger when looking at relatively short time windows in the near future (e.g., 2 sec), and these relations progressively weaken when the length of the sliding window is increased. An important exception was found for the Anterior cingulate cortex that seemed to be the only cluster emerging using longer time – windows (it shows up using 4 seconds during the task and at 8 – 16 seconds

during rest), suggesting that this area could be related to slower (referring to sympathetic) changes in autonomic activity.

The anterior cingulate area is well-known to be involved in top-down modulation of autonomic activity (Craig, 2009; Critchley et al., 2003; Fan et al., 2012; Nagai et al., 2004; Thayer & Lane, 2009). Here we confirmed previous findings on its negatively correlated relation to autonomic fluctuations. The fact that the region was identified also in analyses using longer sliding – window lengths could be explained by ideas presented in Craig’s theory of interoception (Craig, 2009). In this work, Craig shows a model of awareness (defined there as the ability of an organism to experience its own existence) composed by several steps of information integration involving different brain structures, and where the level of integrated material increases step by step. He proposed that posterior insula receives very basic information from the physiological condition of the body, which results in an interoception representation. Other brain structures build on this information, reporting for example conditions of the external environment. The last step is the exchange between previously integrated information (started by posterior insula), motivational, social and cognitive conditions provided by ACC and the anterior insula, which integrates these two pieces of information in order to represent the awareness. In our findings, posterior insular cortex showed significant positive correlation with autonomic activity. The interesting point with respect to Craig’s theory is that we found it only when relating BOLD with short time – windows averaging, suggesting fast, frequently updating changes in the activity of this specific area in the BOLD – AA framework. Anterior cingulate cortex in Craig’s theory acts as the second last character, and this would refer to our finding where ACC shows up later on with respect to posterior insula. Anterior part of the insula in Craig’s theory is the last

cortical structure communicating with anterior cingulate: we also found this insular subregion in our results, only in the steady – state epochs, using a great part of the sliding window lengths. This confirms the fact that this area is generally involved in brain – autonomic exchanges as part of an ongoing information integration cascade.

3.4.2 Differences between resting-state and steady-state task

We identified several areas that showed BOLD-> ANS relations both during rest and during the steady state task: these included posterior insula, posterior cingulate cortex, subcentral gyrus and anterior cingulate cortex. This means that there exist a set of areas generically involved in the BOLD - AA relationship, and we described potential features of this set in the previous section.

We did not find, combining information from all the sliding windows, any area showing relevant effects only during rest, while we found one area, the anterior insula, showing relevant effects only during the task. We did not find any area changing the sign of the relationship given the change in the cognitive environment.

Furthermore, the ACC showed the same BOLD->ANS relation but in different sliding window lengths for task and rest. This suggests that the task context modulates the way in which the ACC determines downstream processes, with a faster regulation rate at task than at rest. Specifically, this area showed negative correlations between its activity and autonomic one when the latter was averaged along 8 and 16 seconds at rest, while during the steady state task it showed the same negative correlation at 4 seconds. This could suggest that when the brain is engaged in a stressful task, the heart starts beating faster (as we also noticed when we described behavioral aspects of heart rate variability during task period), vagal (faster)

influence could start to predominate over the sympathetic one and the frequency of updating activity related to autonomic one in anterior cingulate could increase, causing the cluster to emerge at shorter sliding window lengths.

We found a cluster in the posterior insular region both in task and during rest. It is interesting to evaluate this result in light of Craig's interoception theory (Craig, 2009) where he indicated this area as the starting point of an information integration cascade between external and internal elements. This could suggest that at rest the function of posterior insula is to relate to internal information, without completing the whole integration cascade. Anterior insula cluster showed up only during the task, where it appeared as a general subsystem involved using almost all the sliding window lengths: this could refer to the already known involvement of this area in BOLD – autonomic activity relationship (Thayer & Lane, 2009). The fact that this area here showed up only within the task could suggest that such a high level of information integration is required only when there is an effortful engagement of the participant in something hard to accomplish. An interesting point for future work would be to model the effective connectivity between these two regions, in a fast-TR protocol, during task and rest. It could be expected that there will be more efferent connectivity from posterior to anterior insula, than vice versa

Another interesting point emerging from our results, with respect to studies of the resting state, is the fact that our analysis shows different profiles, with respect to the autonomic activity, for the anterior and posterior cingulate cortices. In most of the studies, so far, these two areas are mentioned together (Beckmann, DeLuca, Devlin, & Smith, 2005; Raichle et al., 2001), as functionally connected by means of a correlation of their time-series. We also already mentioned in the results section that other studies, like the one by Fan and colleagues

(Fan et al., 2012) described a similar dissociation of these two areas with respect to autonomic arousal. There, they reported for the two areas both a similar correlation between BOLD and arousal and the same outcome of a psychophysiological interaction analysis. In our study, on the other hand, we showed that anterior and posterior cingulate cortices are, respectively, negatively and positively correlated with autonomic activity at rest. Furthermore, the two clusters showed up at different sliding window lengths: posterior cingulate using 2 seconds, and anterior using 8 and 16 seconds. Our results suggest an interesting difference in this specific aspect of their function, given that they relate in an opposite way to the autonomic fluctuations. Also, we could suggest different autonomic implications for these two different (one very fast, the other slower) relationships.

3.4.3 Implications for Deactivation theories

We found a relevant set of areas showing a significant relationship between BOLD and autonomic activity and also deactivations. Lin and colleagues (Lin, Hasson, Jovicich, & Robinson, 2011) used calibrated and perfusion fMRI to show how the coupling between blood flow and oxygen consumption in the deactivated region showed similar patterns to activated ones, supporting the argument that deactivations are not occurring for artifactual reasons, but that there exists an organized mode of information processing and neural activity during deactivation. Here we noticed that several areas deactivate during the execution of a cognitive task were actually involved in dealing with autonomic activity during the steady state session. This could suggest that part of their functionality is assigned to the BOLD – autonomic activity relationship and their deactivation during a cognitive task could be an index that they are involved in something different than the actual execution of the task, even if still cognitively relevant, as far as the BOLD – AA is modulated by the task itself.

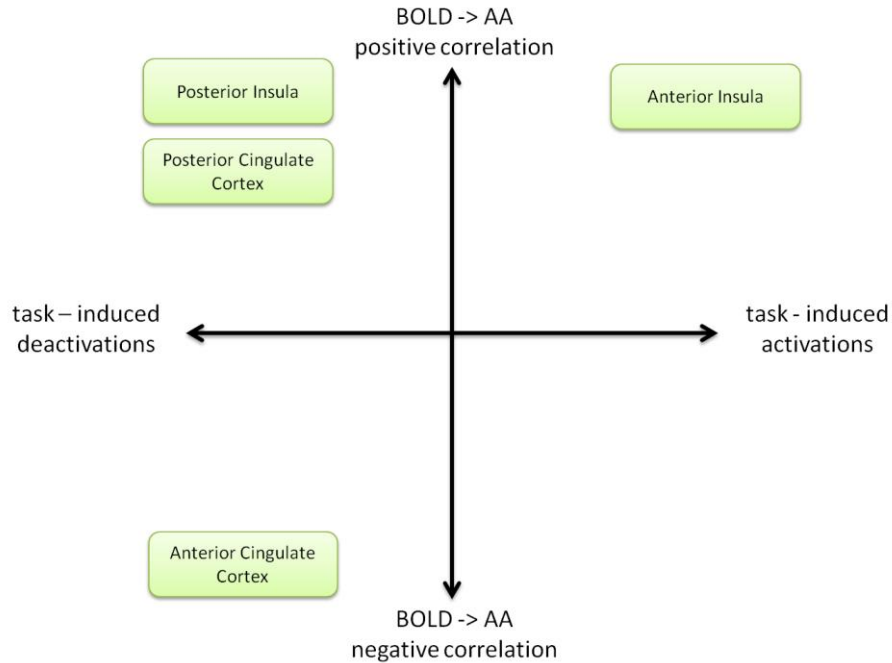


Figure 3.11 – Combination of results from BOLD – AA maps and task – induced effects. Pieces of information coming from the two procedures integrate showing new functional insights about specific areas.

More generally, by combining information coming from BOLD – AA analysis and task induced effects we were able to describe several areas using two distinct parameters, referring to very different aspects of their functionality. For example, we were able to functionally segregate and assign different characteristics to anterior and posterior insula, as it is shown introducing the task – induced effects axis shown in figure 3.11 –. From the same diagram the previously described dichotomy between anterior and posterior cingulate cortex is also noticeable. Combining these two different information could be a way to present and describe the couple brain area – function in a richer way with respect to the plain activation / deactivation sign and amplitude.

3.5 Conclusions

Our work identified a set of areas where BOLD activity systematically precedes autonomic one, including: anterior / posterior cingulate cortex, anterior / posterior insula, anterior / posterior temporal areas and calcarine sulcus. We showed that by filtering in ways that captured different parts of the ANS frequency power, we were able to identify different areas that correlate with those frequency generators of ANS activity.

Importantly, areas showing BOLD->ANS relations were modulated by cognitive environment (i.e.: the engagement in a mental calculation task). Some areas showed such relations only during task performance but not during the resting state by means that several areas show relevant effects with respect to autonomic activity manipulation over different time windows, according to the ongoing (resting or steady-state task) experimental condition.

We combined BOLD – AA information with the one coming out from a functional localizer and we found that several areas significantly related to autonomic activity also showed deactivations, supporting the argument that those are not a plain artifact, but they could be implicated in BOLD – AA relationship. Furthermore, this combination of information describes the functionality of a region in a deeper way, and it could be used to discover features (e.g.: fine tuning functional segregations or showing functional differences avoided by other methods) that were otherwise not observable.

4 General discussion

Summary

In this work we investigated several aspects regarding the relationship between BOLD signal and physiology – related information.

In the first chapter, we reviewed two distinct approaches in considering physiology - related sequences with respect to functional magnetic resonance imaging: one treating physiology – related fluctuations as generators of noise, the other considering them as cognitively relevant.

In the second chapter we considered the first approach - physiology as noise - and proposed a permutation - based, self-consistent benchmark to assess the validity of a physiological noise removal procedure.

In the third chapter, we considered autonomic indicants derived from physiological time - series as meaningful components of the BOLD signal. We described a FMRI experiment building on this, where the goal was to localize brain areas whose activity is directionally related to autonomic one, in a top - down modulation fashion.

Here we review what we concluded in the previous sections and we summarize the general contributions of our findings.

4.1 The problem: two ways to consider physiology – related effects in the BOLD signal

We presented two different views for the role of physiology – related indices in the context of functional magnetic resonance imaging. On one hand, when these indicants are treated purely as physiological noise, the aim of researchers is to remove their effect from the BOLD signal. As far as the co-variates explored in this context are related to purely nuisance variables such as head motion, cardiac-induced tissue motion and changes in CO₂ concentration, this approach is clearly justified. However, the importance of work showing the interesting functional relation between ANS indices and both cortical and subcortical regions should be considered as well. From this perspective, brain regions that play a meaningful functional role in driving or monitoring ANS activity will show activation patterns that correlate with the BOLD signal. Technical procedures that aim to remove the effect of ANS indices from the BOLD signal will also remove, in this context, a meaningful variance component not only when studying processes associated with these regions, but when studying any process in which there might be a confound between cognitive processing and ANS states. Furthermore, as we have outlined, the study of the relation between the ANS and cortical function is an important topic being currently developed and describing the frequencies driving such connectivity, the pathways of information, and the way in which the link between ANS activity and cortical activity can be manipulated are central questions in this domain. (Clearly none of these questions can be answered if ANS indicants were removed from the signal.)

In addition to this, anterior cingulate cortex is a brain hub turning out to be related to both the viewpoints. On one hand, the ACC is a central node in the “Default Mode Network”

(DMN) which is considered to be one of the dominant resting state networks, and studied using functional connectivity methods treating ANS indicants as external factors that represent physiologically-confounding data. On the other hand, the literature we reviewed suggests that the ACC is one of the most important brain mediators between sympathetic information and brain activity, as was shown from studies of clinical and non-clinical populations.

From these considerations, we suggest to carefully evaluate whether a particular research question unambiguously justifies the removal of ANS indicants from the signal.

4.2 Validity of results coming from physiological noise correction

Considering the point of view where physiology – related time – series are treated as generators of noise, we discussed the lack of self – consistency in assessing whether a correction procedure has achieved this aim.

If one aims to rule out the possibility that BOLD synchronization in a certain network simply reflects physiological connectivity, one has to conduct PN-correction and in addition, independently demonstrate that variance related to physiological noise was indeed removed from the dataset. Only at that point is it most meaningful to discuss the features of the PN-corrected data.

The considerations we made and the results we showed could drive to the fact that the problem is not the way in which we look at fMRI data cleaned from physiological effects. We instead think that the real issue is to find a significant chance level to compare the results with. We tried using completely random regressors, but the autocorrelation feature seems to be dramatic when explaining BOLD signal fluctuations. The choice of the permutations

framework seems to be proper for several motivations. First, other participants' physiological data are ready-to-use information that comes from the experimental sessions themselves, with no need of further acquiring. Then, they are on one hand random, because it is unlikely that two subjects, even if they share heart rate and respiration frequencies, report the same trend of these fluctuations in time. On the other hand they still exhibit all the features that these data should have: e.g. autocorrelation and possible artifacts, given that they are taken from the same hardware.

We therefore think that the sampling distribution built from experimental data is an issue that gave us a reliable chance level to compare real data with, maximizing the effects we listed as potentially responsible to loss of validity.

The distance between real-PC result and perm-PC one seems to be a straightforward choice to characterize those datasets where the changes in variance are likely to be related to the physiological time-series. Using this distribution we were able to show, in a reliable way, where those changes could instead reflect computation or other intrinsic features of BOLD signal.

4.3 Physiology as a relevant factor: Directional relationships between BOLD activity and autonomic nervous system fluctuations revealed by fast fMRI acquisition

In the first chapter we reviewed two distinct bodies of literature, and we observed the fact that several areas known to be involved in BOLD – autonomic activity relationship also emerges in functional connectivity studies. In addition to this, using the method described in chapter 2, we found that a relevant amount of variance in the BOLD signal is explained by

physiology – related fluctuations. We ended up with developing a specific study to investigate directional exchanges between BOLD and autonomic activity, given that identifying neural systems related to this issue could potentially contribute to three general questions. First, are there brain areas where cognitive activity is followed by systematic changes in autonomic activity? Second, to the extent that such are found, do they perform similar functions during task and rest states? Finally, are the brain regions showing this relationship brain regions that are considered active by task, or do they also include brain areas deactivated by the task?

To answer these questions in a valid way, we constructed a fast – repetition time (fast-TR) functional magnetic resonance imaging experimental setup, combined with physiological acquisitions, and also included a block-design experiment with a typical TR to quantify whole brain activation and deactivation during the task.

Our work identified a set of areas where BOLD activity systematically precedes autonomic one, including: anterior / posterior cingulate cortex, anterior / posterior insula, anterior / posterior temporal areas and calcarine sulcus. We showed that by filtering in ways that captured different parts of the ANS frequency power, we were able to identify different areas that correlate with those frequency generators of ANS activity.

Importantly, areas showing BOLD->ANS relations were modulated by cognitive environment (i.e.: the engagement in a mental calculation task). Some areas showed such relations only during task performance but not during the resting state by means that several areas show relevant effects with respect to autonomic activity manipulation over different time windows, according to the ongoing (resting or steady-state task) experimental condition.

We combined BOLD – AA information with the one coming out from a functional localizer and we found that several areas significantly related to autonomic activity also showed deactivations, supporting the argument that those are not a plain artifact, but they could be implicated in BOLD – AA relationship. Furthermore, this combination of information describes the functionality of a region in a deeper way, and it could be used to discover features (e.g.: fine tuning functional segregations or showing functional differences avoided by other methods) that were otherwise not observable.

4.4 General contributions

We spent some effort in combining results coming from distinct bodies of work already present in literature. The perspectives discussed in these two collections of studies are commonly not considered together. We observed, indeed, that they share information: for example the fact that anterior cingulate is one of the hubs in the default mode network but also represent one of the most important areas in neurovisceral integration.

We think that bringing together information from these two different fields could help in thinking back at least one issue: the fact that the insertion of cleaning from physiological noise effects procedure has to be carefully calibrated according to the scientific question of the study. Also, if the scientific question somewhat refers to resting state, it could still be worth considering physiological quantities as carriers of cognitively meaningful information instead of noise, given the fact that the function of several brain areas showing relevant activity at rest is also related to autonomic one.

We thought about a self - consistent way to assess the validity of a physiological noise cleaning procedure, and we ended up with proposing and developing a permutation - based

method to quantify if and how the variance removed was actually related to physiology. This way of proceeding is already being implemented in several studies (e.g.: Davis 2012, in preparation) as a quality control step in fMRI preprocessing cascade.

We finally considered the problem of localizing brain regions related with autonomic activity. In doing this, we considered sequences derived from autonomic indicants carrying also behavioral information: this conceptual approach to the physiological time-series was also implemented in other works (Tobia, Iacovella, Davis, & Hasson, 2012), leading to significant results.

Our findings show consistency with both theoretical and experimental studies previously in the topic. With our fast - TR experimental setup and the sliding window - based correlation data analysis method, we found implications relating to generic BOLD - autonomic nervous system relationship, resting state fluctuations and deactivation theories.

Acknowledgements

I would really like to thank Prof. Uri Hasson, for all the important things he taught me in these four years. It has been an extraordinary stroke of luck, for me, to have him assigned as supervisor.

I want to thank all the lab members: Ben, Magda, James, Krishna, Nathan, Mike and Sam, for all the discussions we had during the meetings and in other places.

My former office – mate Michael Tobia is a great person and an outstanding researcher: I learnt a lot from him, and from all the discussion we had in the never-ending office hours.

Gianpaolo De Marchi, Gianpiero Monittola, Mateus Joffily represent a mind – boggling case of scientists able to keep it always informal and informative. I had a great and interesting time with them.

All my fellows in the PhD school, through the years, became more than simple colleagues. I want to mention here the discussions and the great time I had with Silvia, Nicola, Andrea, Francesca, Francesca, Alessio, Gianluca, Luca, Marianna, Laura, Guido, Silvia, Luigi, Lisandro, Michele and Anne.

I thank all the CIMeC staff, in particular the IT guys and Leah because they did an important job in managing several important things.

Finally, I want to spend a word for my former supervisors and colleagues in Rome: Federico Giove, Tommaso Gili and Mauro di Nuzzo. I am happy to see that we were able to keep in touch, and to meet almost every year to discuss and collaborate even now that we are not working together anymore.

5 References

- Bechara, A. (2000). Emotion, Decision Making and the Orbitofrontal Cortex. *Cerebral Cortex*, *10*(3), 295–307. doi:10.1093/cercor/10.3.295
- Beckmann, C. F., DeLuca, M., Devlin, J. T., & Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, *360*(1457), 1001–13. doi:10.1098/rstb.2005.1634
- Berntson, G G, Bigger, J. T., Eckberg, D. L., Grossman, P., Kaufmann, P. G., Malik, M., Nagaraja, H. N., et al. (1997). Heart rate variability: origins, methods, and interpretive caveats. *Psychophysiology*, *34*(6), 623–48. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9401419>
- Berntson, G G, Cacioppo, J. T., & Quigley, K. S. (1991). Autonomic determinism: the modes of autonomic control, the doctrine of autonomic space, and the laws of autonomic constraint. *Psychological review*, *98*(4), 459–87. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/1660159>
- Berntson, Gary G, Lozano, D. L., & Chen, Y.-J. (2005). Filter properties of root mean square successive difference (RMSSD) for heart rate. *Psychophysiology*, *42*(2), 246–52. doi:10.1111/j.1469-8986.2005.00277.x
- Bianciardi, M., Fukunaga, M., van Gelderen, P., Horovitz, S. G., de Zwart, J. A., Shmueli, K., & Duyn, J. H. (2009). Sources of functional magnetic resonance imaging signal fluctuations in the human brain at rest: a 7 T study. *Magnetic resonance imaging*, *27*(8), 1019–29. doi:10.1016/j.mri.2009.02.004
- Bianciardi, M., van Gelderen, P., Duyn, J. H., Fukunaga, M., & de Zwart, J. A. (2009). Making the most of fMRI at 7 T by suppressing spontaneous signal fluctuations. *NeuroImage*, *44*(2), 448–54. doi:10.1016/j.neuroimage.2008.08.037
- Birn, R. M. (2012). The role of physiological noise in resting-state functional connectivity. *NeuroImage*, *62*(2), 864–70. doi:10.1016/j.neuroimage.2012.01.016
- Birn, R. M., Diamond, J. B., Smith, M. A., & Bandettini, P. A. (2006). Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. *NeuroImage*, *31*(4), 1536–48. doi:10.1016/j.neuroimage.2006.02.048
- Birn, R. M., Smith, M. A., Jones, T. B., & Bandettini, P. A. (2008). The respiration response function: the temporal dynamics of fMRI signal fluctuations related to changes in respiration. *NeuroImage*, *40*(2), 644–54. doi:10.1016/j.neuroimage.2007.11.059

- Biswal, B., Zerrin Yetkin, F., Haughton, V. M., & Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine*, *34*(4), 537–541. doi:10.1002/mrm.1910340409
- Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network: anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, *1124*, 1–38. doi:10.1196/annals.1440.011
- Buxton, R. B., Uludağ, K., Dubowitz, D. J., & Liu, T. T. (2004). Modeling the hemodynamic response to brain activation. *NeuroImage*, *23 Suppl 1*, S220–33. doi:10.1016/j.neuroimage.2004.07.013
- Buxton, R. B., Wong, E. C., & Frank, L. R. (1998). Dynamics of blood flow and oxygenation changes during brain activation: The balloon model. *Magnetic Resonance in Medicine*, *39*(6), 855–864. doi:10.1002/mrm.1910390602
- Cannon, W. B. (1987). The James-Lange theory of emotions: a critical examination and an alternative theory. By Walter B. Cannon, 1927. *The American journal of psychology*, *100*(3-4), 567–86. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/3322057>
- Chang, C., Cunningham, J. P., & Glover, G. H. (2009). Influence of heart rate on the BOLD signal: the cardiac response function. *NeuroImage*, *44*(3), 857–69. doi:10.1016/j.neuroimage.2008.09.029
- Chang, C., & Glover, G. H. (2009). Effects of model-based physiological noise correction on default mode network anti-correlations and correlations. *NeuroImage*, *47*(4), 1448–59. doi:10.1016/j.neuroimage.2009.05.012
- Chen, C.-C., Tyler, C. W., & Baseler, H. A. (2003). Statistical properties of BOLD magnetic resonance activity in the human brain. *NeuroImage*, *20*(2), 1096–109. doi:10.1016/S1053-8119(03)00358-6
- Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and biomedical research, an international journal*, *29*(3), 162–73. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/8812068>
- Craig, A. (2002). How do you feel? Interoception: the sense of the physiological condition of the body, *3*(August).
- Craig, A. (2009). How do you feel--now? The anterior insula and human awareness. *Nature reviews. Neuroscience*, *10*(1), 59–70. doi:10.1038/nrn2555
- Critchley, H D. (2005). Neural mechanisms of autonomic, affective, and cognitive integration. *The Journal of comparative neurology*, *493*(1), 154–66. doi:10.1002/cne.20749

- Critchley, H D, Corfield, D. R., Chandler, M. P., Mathias, C. J., & Dolan, R. J. (2000). Cerebral correlates of autonomic cardiovascular arousal: a functional neuroimaging investigation in humans. *The Journal of physiology*, *523 Pt 1*, 259–70. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2269796&tool=pmcentrez&rendertype=abstract>
- Critchley, Hugo D, Mathias, C. J., Josephs, O., O’Doherty, J., Zanini, S., Dewar, B.-K., Cipolotti, L., et al. (2003). Human cingulate cortex and autonomic control: converging neuroimaging and clinical evidence. *Brain: a journal of neurology*, *126*(Pt 10), 2139–52. doi:10.1093/brain/awg216
- Dagli, M. S., Ingeholm, J. E., & Haxby, J. V. (1999). Localization of cardiac-induced signal change in fMRI. *NeuroImage*, *9*(4), 407–15. doi:10.1006/nimg.1998.0424
- Davis, T. L., Kwong, K. K., Weisskoff, R. M., & Rosen, B. R. (1998). Calibrated functional MRI: Mapping the dynamics of oxidative metabolism. *Proceedings of the National Academy of Sciences*, *95*(4), 1834–1839. doi:10.1073/pnas.95.4.1834
- De Luca, M., Beckmann, C. F., De Stefano, N., Matthews, P. M., & Smith, S. M. (2006). fMRI resting state networks define distinct modes of long-distance interactions in the human brain. *NeuroImage*, *29*(4), 1359–67. doi:10.1016/j.neuroimage.2005.08.035
- De Neys, W., Moyens, E., & Vansteenwegen, D. (2010). Feeling we’re biased: autonomic arousal and reasoning conflict. *Cognitive, affective & behavioral neuroscience*, *10*(2), 208–16. doi:10.3758/CABN.10.2.208
- Deen, B., Pitskel, N. B., & Pelphrey, K. a. (2011). Three systems of insular functional connectivity identified with cluster analysis. *Cerebral cortex (New York, N.Y. □: 1991)*, *21*(7), 1498–506. doi:10.1093/cercor/bhq186
- Duff, E. P., Johnston, L. A., Xiong, J., Fox, P. T., Mareels, I., & Egan, G. F. (2008). The power of spectral density analysis for mapping endogenous BOLD signal fluctuations. *Human brain mapping*, *29*(7), 778–90. doi:10.1002/hbm.20601
- Elghozi, J.-L., & Julien, C. (2007). Sympathetic control of short-term heart rate variability and its pharmacological modulation. *Fundamental & clinical pharmacology*, *21*(4), 337–47. doi:10.1111/j.1472-8206.2007.00502.x
- Evans, K. C. (2010). Cortico-limbic circuitry and the airways: insights from functional neuroimaging of respiratory afferents and efferents. *Biological psychology*, *84*(1), 13–25. doi:10.1016/j.biopsycho.2010.02.005
- Fan, J., Xu, P., Van Dam, N. T., Eilam-Stock, T., Gu, X., Luo, Y.-J., & Hof, P. R. (2012). Spontaneous brain activity relates to autonomic arousal. *The Journal of neuroscience □: the official journal of the Society for Neuroscience*, *32*(33), 11176–86. doi:10.1523/JNEUROSCI.1172-12.2012

- Forman, S. D., Cohen, J. D., Fitzgerald, M., Eddy, W. F., Mintun, M. A., & Noll, D. C. (1995). Improved Assessment of Significant Activation in Functional Magnetic Resonance Imaging (fMRI): Use of a Cluster-Size Threshold. *Magnetic Resonance in Medicine*, 33(5), 636–647. doi:10.1002/mrm.1910330508
- Fox, M. D., Zhang, D., Snyder, A. Z., & Raichle, M. E. (2009). The global signal and observed anticorrelated resting state brain networks. *Journal of neurophysiology*, 101(6), 3270–83. doi:10.1152/jn.90777.2008
- Fox, P. T. (1986). Focal Physiological Uncoupling of Cerebral Blood Flow and Oxidative Metabolism during Somatosensory Stimulation in Human Subjects. *Proceedings of the National Academy of Sciences*, 83(4), 1140–1144. doi:10.1073/pnas.83.4.1140
- Gianaros, P. J., Van Der Veen, F. M., & Jennings, J. R. (2004). Regional cerebral blood flow correlates with heart period and high-frequency heart period variability during working-memory tasks: Implications for the cortical and subcortical regulation of cardiac autonomic activity. *Psychophysiology*, 41(4), 521–30. doi:10.1111/1469-8986.2004.00179.x
- Giardino, N. D., Friedman, S. D., & Dager, S. R. (2007). Anxiety, respiration, and cerebral blood flow: implications for functional brain imaging. *Comprehensive psychiatry*, 48(2), 103–12. doi:10.1016/j.comppsy.2006.11.001
- Giove, F., Gili, T., Iacovella, V., Macaluso, E., & Maraviglia, B. (2009). Images-based suppression of unwanted global signals in resting-state functional connectivity studies. *Magnetic resonance imaging*, 27(8), 1058–64. doi:10.1016/j.mri.2009.06.004
- Glover, G. H., Li, T. Q., & Ress, D. (2000). Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magnetic resonance in medicine*: official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine, 44(1), 162–7. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10893535>
- Gray, M. A., Minati, L., Harrison, N. A., Gianaros, P. J., Napadow, V., & Critchley, H. D. (2009). Physiological recordings: basic concepts and implementation during functional magnetic resonance imaging. *NeuroImage*, 47(3), 1105–15. doi:10.1016/j.neuroimage.2009.05.033
- Gray, M. A., Rylander, K., Harrison, N. A., Wallin, B. G., & Critchley, H. D. (2009). Following one's heart: cardiac rhythms gate central initiation of sympathetic reflexes. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 29(6), 1817–25. doi:10.1523/JNEUROSCI.3363-08.2009
- Greicius, M. D., & Menon, V. (2004). Default-mode activity during a passive sensory task: uncoupled from deactivation but impacting activation. *Journal of cognitive neuroscience*, 16(9), 1484–92. doi:10.1162/0898929042568532

- Guimarães, H. N., & Santos, R. a. (1998). A comparative analysis of preprocessing techniques of cardiac event series for the study of heart rhythm variability using simulated signals. *Brazilian journal of medical and biological research = Revista brasileira de pesquisas médicas e biológicas / Sociedade Brasileira de Biofísica ... [et al.]*, 31(3), 421–30. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9698793>
- Harrison, N. A., Gray, M. A., Gianaros, P. J., & Critchley, H. D. (2010). The embodiment of emotional feelings in the brain. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 30(38), 12878–84. doi:10.1523/JNEUROSCI.1725-10.2010
- Harvey, A. K., Pattinson, K. T. S., Brooks, J. C. W., Mayhew, S. D., Jenkinson, M., & Wise, R. G. (2008). Brainstem functional magnetic resonance imaging: disentangling signal from physiological noise. *Journal of magnetic resonance imaging: JMRI*, 28(6), 1337–44. doi:10.1002/jmri.21623
- Hu, J., Lee, J.-M., Gao, J., White, K. D., & Crosson, B. (2008). Assessing a signal model and identifying brain activity from fMRI data by a detrending-based fractal analysis. *Brain structure & function*, 212(5), 417–26. doi:10.1007/s00429-007-0166-9
- Hyder, F., & Rothman, D. L. (2010). Neuronal correlate of BOLD signal fluctuations at rest: err on the side of the baseline. *Proceedings of the National Academy of Sciences of the United States of America*, 107(24), 10773–4. doi:10.1073/pnas.1005135107
- Jo, H. J., Saad, Z. S., Simmons, W. K., Milbury, L. A., & Cox, R. W. (2010). Mapping sources of correlation in resting state FMRI, with artifact detection and removal. *NeuroImage*, 52(2), 571–82. doi:10.1016/j.neuroimage.2010.04.246
- Johnsen, E. L., Tranel, D., Lutgendorf, S., & Adolphs, R. (2009). A neuroanatomical dissociation for emotion induced by music. *International journal of psychophysiology: official journal of the International Organization of Psychophysiology*, 72(1), 24–33. doi:10.1016/j.ijpsycho.2008.03.011
- Jones, T. B., Bandettini, P. A., & Birn, R. M. (2008). Integration of motion correction and physiological noise regression in fMRI. *NeuroImage*, 42(2), 582–90. doi:10.1016/j.neuroimage.2008.05.019
- Kastrup, A., Li, T. Q., Glover, G. H., & Moseley, M. E. (1999). Cerebral blood flow-related signal changes during breath-holding. *AJNR. American journal of neuroradiology*, 20(7), 1233–8. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10472977>
- Katura, T., Tanaka, N., Obata, A., Sato, H., & Maki, A. (2006). Quantitative evaluation of interrelations between spontaneous low-frequency oscillations in cerebral hemodynamics and systemic cardiovascular dynamics. *NeuroImage*, 31(4), 1592–600. doi:10.1016/j.neuroimage.2006.02.010

- Kawada, T., Ikeda, Y., Sugimachi, M., Shishido, T., Kawaguchi, O., Yamazaki, T., Alexander, J., et al. (1996). Bidirectional augmentation of heart rate regulation by autonomic nervous system in rabbits. *The American journal of physiology*, *271*(1 Pt 2), H288–95. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/8760187>
- Kawashima, R., Taira, M., Okita, K., Inoue, K., Tajima, N., Yoshida, H., Sasaki, T., et al. (2004). A functional MRI study of simple arithmetic—a comparison between children and adults. *Cognitive Brain Research*, *18*(3), 227–233. doi:10.1016/j.cogbrainres.2003.10.009
- Kobayashi, N., Yoshino, A., Takahashi, Y., & Nomura, S. (2007). Autonomic arousal in cognitive conflict resolution. *Autonomic neuroscience: basic & clinical*, *132*(1-2), 70–5. doi:10.1016/j.autneu.2006.09.004
- Lang, P. J. (1994). The varieties of emotional experience: a meditation on James-Lange theory. *Psychological review*, *101*(2), 211–21. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/8022956>
- Lin, P., Hasson, U., Jovicich, J., & Robinson, S. (2011). A neuronal basis for task-negative responses in the human brain. *Cerebral cortex (New York, N.Y. □: 1991)*, *21*(4), 821–30. doi:10.1093/cercor/bhq151
- Logothetis, N K, Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, *412*(6843), 150–7. doi:10.1038/35084005
- Logothetis, Nikos K. (2008). What we can do and what we cannot do with fMRI. *Nature*, *453*(7197), 869–78. doi:10.1038/nature06976
- Logothetis, Nikos K. (2003). The Underpinnings of the BOLD Functional Magnetic Resonance Imaging Signal. *J. Neurosci.*, *23*(10), 3963–3971. Retrieved from <http://www.jneurosci.org/content/23/10/3963.full>
- Lu, H., Zhao, C., Ge, Y., & Lewis-Amezcu, K. (2008). Baseline blood oxygenation modulates response amplitude: Physiologic basis for intersubject variations in functional MRI signals. *Magnetic resonance in medicine: official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, *60*(2), 364–72. doi:10.1002/mrm.21686
- Lund, T. E., Madsen, K. H., Sidaros, K., Luo, W.-L., & Nichols, T. E. (2006). Non-white noise in fMRI: does modelling have an impact? *NeuroImage*, *29*(1), 54–66. doi:10.1016/j.neuroimage.2005.07.005
- Macey, P. M., Macey, K. E., Kumar, R., & Harper, R. M. (2004). A method for removal of global effects from fMRI time series. *NeuroImage*, *22*(1), 360–6. doi:10.1016/j.neuroimage.2003.12.042

- Malik, M. (1996). Heart Rate Variability: Standards of Measurement, Physiological Interpretation, and Clinical Use. *Circulation*, 93(5), 1043–1065. doi:10.1161/01.CIR.93.5.1043
- McNames, J., Thong, T., & Goldstein, B. (2003). Reliability and accuracy of heart rate variability metrics versus ECG segment duration. *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE Cat. No.03CH37439)* (pp. 212–215). IEEE. doi:10.1109/IEMBS.2003.1279574
- Mezer, A., Yovel, Y., Pasternak, O., Gorfine, T., & Assaf, Y. (2009). Cluster analysis of resting-state fMRI time series. *NeuroImage*, 45(4), 1117–25. doi:10.1016/j.neuroimage.2008.12.015
- Montano, N., Porta, A., & Malliani, A. (2001). Evidence for central organization of cardiovascular rhythms. *Annals of the New York Academy of Sciences*, 940, 299–306. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11458687>
- Montano, N., Ruscone, T. G., Porta, a., Lombardi, F., Pagani, M., & Malliani, a. (1994). Power spectrum analysis of heart rate variability to assess the changes in sympathovagal balance during graded orthostatic tilt. *Circulation*, 90(4), 1826–1831. doi:10.1161/01.CIR.90.4.1826
- Murphy, K., Birn, R. M., Handwerker, D. A., Jones, T. B., & Bandettini, P. A. (2009). The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? *NeuroImage*, 44(3), 893–905. doi:10.1016/j.neuroimage.2008.09.036
- Murphy, K., Harris, A. D., & Wise, R. G. (2011). Robustly measuring vascular reactivity differences with breath-hold: normalising stimulus-evoked and resting state BOLD fMRI data. *NeuroImage*, 54(1), 369–79. doi:10.1016/j.neuroimage.2010.07.059
- Nagai, Y., Critchley, H. D., Featherstone, E., Trimble, M. R., & Dolan, R. J. (2004). Activity in ventromedial prefrontal cortex covaries with sympathetic skin conductance level: a physiological account of a “default mode” of brain function. *NeuroImage*, 22(1), 243–51. doi:10.1016/j.neuroimage.2004.01.019
- Noll, D. C., & Schneider, W. (n.d.). Theory, simulation, and compensation of physiological motion artifacts in functional MRI. *Proceedings of 1st International Conference on Image Processing* (Vol. 3, pp. 40–44). IEEE Comput. Soc. Press. doi:10.1109/ICIP.1994.413892
- Nussinovitch, U., Elishkevitz, K. P., Katz, K., Nussinovitch, M., Segev, S., Volovitz, B., & Nussinovitch, N. (2011). Reliability of Ultra-Short ECG Indices for Heart Rate Variability. *Annals of noninvasive electrocardiology: the official journal of the International Society for Holter and Noninvasive Electrocardiology, Inc*, 16(2), 117–22. doi:10.1111/j.1542-474X.2011.00417.x

- Obrig, H., Neufang, M., Wenzel, R., Kohl, M., Steinbrink, J., Einhüpl, K., & Villringer, A. (2000). Spontaneous low frequency oscillations of cerebral hemodynamics and metabolism in human adults. *NeuroImage*, *12*(6), 623–39. doi:10.1006/ning.2000.0657
- Pollatos, O., Schandry, R., Auer, D. P., & Kaufmann, C. (2007). Brain structures mediating cardiovascular arousal and interoceptive awareness. *Brain research*, *1141*, 178–87. doi:10.1016/j.brainres.2007.01.026
- Raichle, M. E., MacLeod, a M., Snyder, a Z., Powers, W. J., Gusnard, D. a, & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(2), 676–82. doi:10.1073/pnas.98.2.676
- Renvall, V., & Hari, R. (2009). Transients may occur in functional magnetic resonance imaging without physiological basis. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(48), 20510–4. doi:10.1073/pnas.0911265106
- Saad, Z. S., Glen, D. R., Chen, G., Beauchamp, M. S., Desai, R., & Cox, R. W. (2009). A new method for improving functional-to-structural MRI alignment using local Pearson correlation. *NeuroImage*, *44*(3), 839–48. doi:10.1016/j.neuroimage.2008.09.037
- Schroeder, E. B., Whitsel, E. A., Evans, G. W., Prineas, R. J., Chambless, L. E., & Heiss, G. (2004). Repeatability of heart rate variability measures. *Journal of electrocardiology*, *37*(3), 163–72. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15286929>
- Schölvinck, M. L., Maier, A., Ye, F. Q., Duyn, J. H., & Leopold, D. A. (2010). Neural basis of global resting-state fMRI activity. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(22), 10238–43. doi:10.1073/pnas.0913110107
- Shmuel, A., Augath, M., Oeltermann, A., & Logothetis, N. K. (2006). Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. *Nature neuroscience*, *9*(4), 569–77. doi:10.1038/nn1675
- Shmueli, K., van Gelderen, P., de Zwart, J. A., Horovitz, S. G., Fukunaga, M., Jansma, J. M., & Duyn, J. H. (2007). Low-frequency fluctuations in the cardiac rate as a source of variance in the resting-state fMRI BOLD signal. *NeuroImage*, *38*(2), 306–20. doi:10.1016/j.neuroimage.2007.07.037
- Soldati, N., Robinson, S., Persello, C., Jovicich, J., & Bruzzone, L. (2009). Automatic classification of brain resting states using fMRI temporal signals. *Electronics Letters*, *45*(1), 19. doi:10.1049/el:20092178
- Teichert, T., Grinband, J., Hirsch, J., & Ferrera, V. P. (2010). Effects of heartbeat and respiration on macaque fMRI: implications for functional connectivity. *Neuropsychologia*, *48*(7), 1886–94. doi:10.1016/j.neuropsychologia.2009.11.026

- Thayer, J. F., Ahs, F., Fredrikson, M., Sollers, J. J., & Wager, T. D. (2012). A meta-analysis of heart rate variability and neuroimaging studies: implications for heart rate variability as a marker of stress and health. *Neuroscience and biobehavioral reviews*, *36*(2), 747–56. doi:10.1016/j.neubiorev.2011.11.009
- Thayer, J. F., & Lane, R. D. (2009). Claude Bernard and the heart-brain connection: further elaboration of a model of neurovisceral integration. *Neuroscience and biobehavioral reviews*, *33*(2), 81–8. doi:10.1016/j.neubiorev.2008.08.004
- Thayer, J. F., & Sternberg, E. (2006). Beyond Heart Rate Variability Vagal Regulation of Allostatic Systems, *372*, 361–372. doi:10.1196/annals.1366.014
- Thomason, M. E., Burrows, B. E., Gabrieli, J. D. E., & Glover, G. H. (2005). Breath holding reveals differences in fMRI BOLD signal in children and adults. *NeuroImage*, *25*(3), 824–37. doi:10.1016/j.neuroimage.2004.12.026
- Tobia, M. J., Iacovella, V., Davis, B., & Hasson, U. (2012). Neural systems mediating recognition of changes in statistical regularities. *NeuroImage*, *63*(3), 1730–42. doi:10.1016/j.neuroimage.2012.08.017
- Triantafyllou, C., Hoge, R. D., Krueger, G., Wiggins, C. J., Potthast, a, Wiggins, G. C., & Wald, L. L. (2005). Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *NeuroImage*, *26*(1), 243–50. doi:10.1016/j.neuroimage.2005.01.007
- Van de Moortele, P.-F., Pfeuffer, J., Glover, G. H., Ugurbil, K., & Hu, X. (2002). Respiration-induced B0 fluctuations and their spatial distribution in the human brain at 7 Tesla. *Magnetic resonance in medicine* □: official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine, *47*(5), 888–95. doi:10.1002/mrm.10145
- Wager, T. D., Waugh, C. E., Lindquist, M., Noll, D. C., Fredrickson, B. L., & Taylor, S. F. (2009). Brain mediators of cardiovascular responses to social threat: part I: Reciprocal dorsal and ventral sub-regions of the medial prefrontal cortex and heart-rate reactivity. *NeuroImage*, *47*(3), 821–35. doi:10.1016/j.neuroimage.2009.05.043
- Wager, T. D., van Ast, V. A., Hughes, B. L., Davidson, M. L., Lindquist, M. A., & Ochsner, K. N. (2009). Brain mediators of cardiovascular responses to social threat, part II: Prefrontal-subcortical pathways and relationship with anxiety. *NeuroImage*, *47*(3), 836–51. doi:10.1016/j.neuroimage.2009.05.044
- Wise, R. G., Ide, K., Poulin, M. J., & Tracey, I. (2004). Resting fluctuations in arterial carbon dioxide induce significant low frequency variations in BOLD signal. *NeuroImage*, *21*(4), 1652–64. doi:10.1016/j.neuroimage.2003.11.025

- Wu, C. W., Gu, H., Lu, H., Stein, E. A., Chen, J.-H., & Yang, Y. (2008). Frequency specificity of functional connectivity in brain networks. *NeuroImage*, *42*(3), 1047–55. doi:10.1016/j.neuroimage.2008.05.035
- van Buuren, M., Gladwin, T. E., Zandbelt, B. B., van den Heuvel, M., Ramsey, N. F., Kahn, R. S., & Vink, M. (2009). Cardiorespiratory effects on default-mode network activity as measured with fMRI. *Human brain mapping*, *30*(9), 3031–42. doi:10.1002/hbm.20729