

Relationships between correlated spikes, oxygen and LFP in the resting-state primate

Jingfeng M. Li^a, Benjamin T. Acland^a, Alexander S. Brenner^b, William J. Bentley^a, Lawrence H. Snyder^{a,b,*}

^a Department of Neuroscience, Washington University School of Medicine, 660 S Euclid Ave, Box 8108, St Louis, MO 63110, USA

^b Department of Biomedical Engineering, Washington University, St Louis, MO 63130, USA

ARTICLE INFO

Keywords:

Default mode network
Functional connectivity
Magnetic resonance imaging
Neurohemodynamic coupling
Oxygen polarography

ABSTRACT

Resting-state functional MRI (rsfMRI) provides a view of human brain organization based on correlation patterns of blood oxygen level dependent (BOLD) signals recorded across the whole brain. The neural basis of resting-state BOLD fluctuations and their correlation remains poorly understood. We simultaneously recorded oxygen level, spikes, and local field potential (LFP) at multiple sites in awake, resting monkeys. Following a spike, the average local oxygen and LFP voltage responses each resemble a task-driven BOLD response, with LFP preceding oxygen by 0.5 s. Between sites, features of the long-range correlation patterns of oxygen, LFP, and spikes are similar to features seen in rsfMRI. Most of the variance shared between sites lies in the infraslow frequency band (0.01–0.1 Hz) and in the infraslow envelope of higher-frequency bands (e.g. gamma LFP). While gamma LFP and infraslow LFP are both strong correlates of local oxygen, infraslow LFP explains significantly more of the variance shared between correlated oxygen signals than any other electrophysiological signal. Together these findings are consistent with a causal relationship between infraslow LFP and long-range oxygen correlations in the resting state.

1. Introduction

BOLD signals from resting humans, non-human primates, and rodents exhibit infraslow (0.01 - 0.1 Hz) fluctuations that are highly correlated among spatially distant regions. The pattern of these correlations contains reproducible inter-regional groupings, or networks, that tend to co-fluctuate during tasks (Fox et al., 2005; Gorges et al., 2017; Ma et al., 2016; Raichle et al., 2001). Resting-state BOLD network structure shows subtle, apparently functionally significant, variation among individuals and between control and patient populations (Peer et al., 2017). However, resting-state BOLD networks do not strictly correspond to any known pattern of anatomical connectivity between neurons, and their neural bases and functional roles (if any) remain unclear.

Much evidence indicates that infraslow BOLD signal correlations comprising resting-state networks reflect correlated infraslow fluctuations in neuronal activity. Consistent relationships have been found between local fluctuations in at-rest BOLD and spikes (Magri et al., 2012; Shmuel and Leopold, 2008), band-limited LFP power (Magri et al., 2012; Pan et al., 2011; Shmuel and Leopold, 2008; Thompson et al., 2013) and infraslow LFP (0.01 - 0.1 Hz) (Pan et al., 2013). Nir et al. (2008) demonstrate long-range correlations in neuronal spiking that occur at roughly the same time scale as resting-state BOLD correlations (Nir et al., 2008).

Resting-state correlations also exist in band-limited LFP power and in infraslow LFP (He et al., 2008; Li et al., 2014). The spatial structures of these electrophysiological correlations match those of BOLD correlations. The spatial match to BOLD correlations is particularly good for infraslow LFP (He et al., 2008), suggesting that the processes underlying infraslow LFP may be closely related to those underlying correlated BOLD signals (Khader et al., 2008). However there has been no direct demonstration in a primate model showing a consistent relationship between, on the one hand, correlations in BOLD recorded at two locations, and on the other hand, correlations in any electrophysiological signals recorded at those same two locations. This gap is an important one to fill, since infraslow LFP and LFP power could be driven, in whole or in part, by non-neural elements such as glia or by hemodynamic processes (for review see Khader et al., 2008) and BOLD signals are also affected by non-neuronal sources (Birn, 2012; Schulz et al., 2012; Tong et al., 2015; Wang et al., 2018). Thus, it is possible that long-range correlations in the different signal types could be driven by different sources.

To definitively establish the temporal relationship between correlated oxygen and correlated electrical activity, the two signals must be recorded simultaneously. To our knowledge, this has only been done in mice (He et al., 2018; Kozberg and Hillman, 2016; Murphy et al., 2018; Vazquez et al., 2014). Using a system to simultaneously record

* Corresponding author at: Department of Neuroscience, Washington University School of Medicine, 660 S Euclid Ave, Box 8108, St Louis, MO 63110, USA.
E-mail address: larry@eye-hand.wustl.edu (L.H. Snyder).

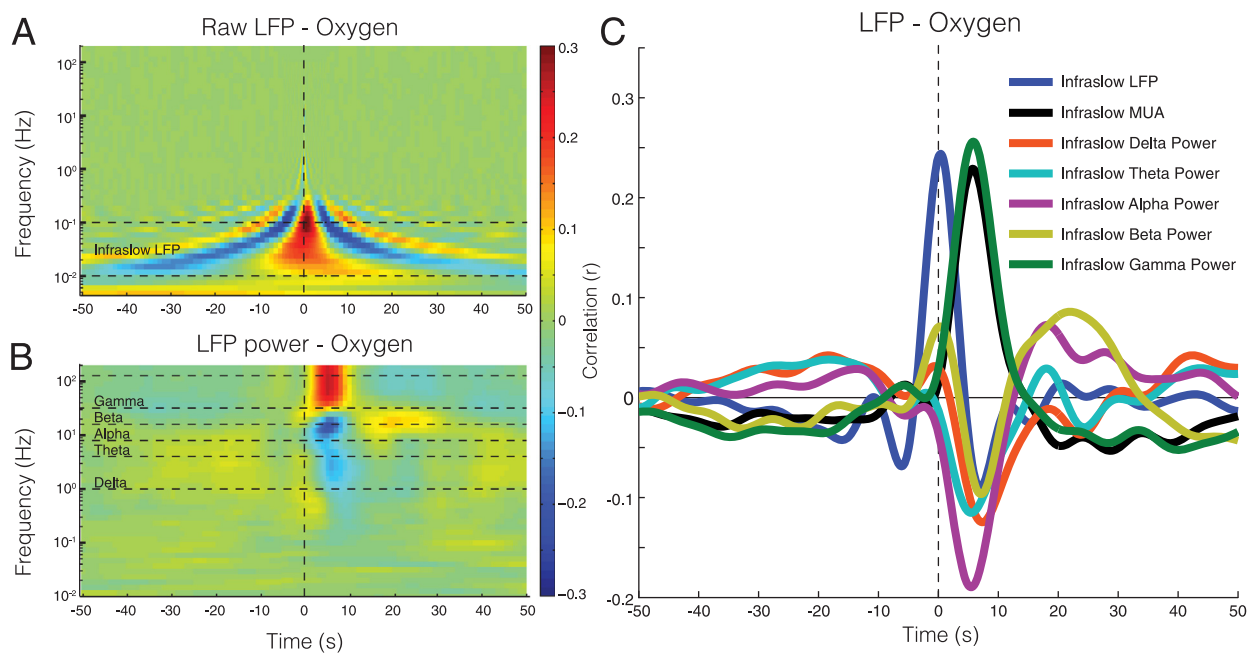


Fig 1. Lagged linear correlation between oxygen and electrophysiological signals recorded from the same site. **A)** Lagged correlation between oxygen and the raw LFP signal recorded in IPS (For V3 and PCC, see Fig. S6). Both the raw LFP signal and the oxygen signal were split into many frequency bands, and then correlations were calculated band by band. The X-axis denotes lag, the Y-axis denotes the frequency band and the color scale denotes correlation strength. Positive lags mean that electrophysiological signals were shifted forwards in time before computing the correlation. Correlations were most prominent at infralow frequencies (0.01 to 0.1 Hz) with little or no lag. **B)** Lagged correlation between oxygen and LFP power. Format is the same as A, except that once LFP power was computed, both the LFP power signal and the oxygen signal were filtered at 0.01 to 0.1 Hz, prior to computing the correlation. Correlations were most prominent for gamma band power, with a lag of ~6 s. **C)** Lagged correlation between infralow oxygen and infralow LFP voltage (see panel A), infralow band-limited LFP power within standard EEG bands (see panel B), and infralow MUA. Infralow LFP, gamma power and MUA have the strongest correlations with oxygen, significantly stronger than the next-highest correlated signal (Infralow LFP vs. beta LFP: $p < 0.05$, $t(58) = 2.414$; gamma LFP vs. beta LFP: $p < 0.05$, $t(58) = 3.378$), but not significantly different from one another ($p = 0.8$, $t(58) = 0.2121$). Peak correlation for infralow LFP is at 0.5 s, which is much shorter than that for gamma power (5.8 s).

oxygen levels and electrophysiology in non-human primates that captures slow, long-range oxygen correlations similar to those seen with BOLD fMRI (Li et al., 2015), we demonstrate that similar correlations exist in the firing rate of extracellular action potentials (spiking activity), that spikes in one hemisphere are correlated with oxygen levels in homotopic cortex in the other hemisphere, and that infralow LFP is the strongest electrophysiological predictor of cross-hemisphere oxygen correlations.

2. Results

2.1. Lagged correlation and single spike responses of oxygen and infralow LFP

We simultaneously recorded from four sites in two different functional networks in awake monkeys resting quietly in a dark room and examined correlations between signals (Supplemental Fig. S15). Two of the recording sites were in left and right posterior cingulate cortex (PCC), which is part of the task-negative default mode network in macaques. (Vincent et al., 2007). The other sites were either in left and right V3 (visual network; monkeys P and L) or left and right intraparietal sulcus (IPS) (attention network; monkey E). In macaques, these areas participate in task-positive networks functionally involved in visual processing, oculomotor control, and attention (Babapoor-Farrokhran et al., 2013; Felleman et al., 1997; Hutchison et al., 2012; Mars et al., 2011; Vincent et al., 2007). Single-unit spiking activity, multi-unit activity (MUA), LFP, and oxygen signals were recorded simultaneously from each site. We first examined the local relationships between infralow oxygen and an array of electrophysiological signals. From LFP, we extracted both the band-limited power of LFP signals and the infralow (0.01–0.1 Hz) fluctuations of raw LFP voltage.

To examine the temporal relationships between signals, we identified the direction and amount each electrophysiological signal had to be shifted (in time) to maximize the absolute value of its cross-correlation with oxygen. Results for most signals were similar across regions, but are shown separately in Fig. 1 (IPS) and Fig. S6 (V3 and PCC). Every tested electrophysiological signal leads oxygen (Fig. 1). Fig. 1A shows the lagged correlation between infralow (0.01–0.1 Hz) oxygen and raw LFP signals at different frequencies. The correlation is prominent in low frequencies (below 0.3 Hz; see also Supplemental Fig. S4), peaking at 0.05 Hz with LFP leading (preceding) oxygen by 0.5 s. Fig. 1B shows the lagged correlation between infralow fluctuations of band-limited LFP power and infralow oxygen signals. This analysis reveals strong lagged correlations above 30 Hz, peaking at 76 Hz with LFP power leading oxygen by 6 s. Fig. 1C shows cross-correlations between infralow oxygen and conventional LFP power bands, infralow LFP and MUA. Of these, oxygen most strongly correlates with infralow LFP, gamma power and MUA across all recorded regions (Fig. S6). Infralow LFP and gamma power have slightly higher correlation coefficients than MUA. Gamma power and MUA lead (precede) oxygen by 6 s. Beta, alpha and theta power show lower peak correlations at longer leads (8–20 s) and with considerable interregional differences (Fig. S6). Previous studies have also reported that electrophysiological signals lead oxygen, although the lead varies across studies (3–8 s) as does the particular frequency band with the strongest correlation, likely due to differences in species or behavioral state (Magri et al., 2012; Murayama et al., 2010; Pan et al., 2011; Schölvinck et al., 2010; Shmuel and Leopold, 2008). Infralow LFP leads oxygen by only 0.5 s, much less than any other electrophysiological signal. (This is also much shorter than the LFP lead found by Pan et al. (2011) [2.5 s under DMED, and 4 s under ISO], perhaps reflecting a difference between awake monkeys and anesthetized rats.) To summarize these results, we see that there is coupling between a

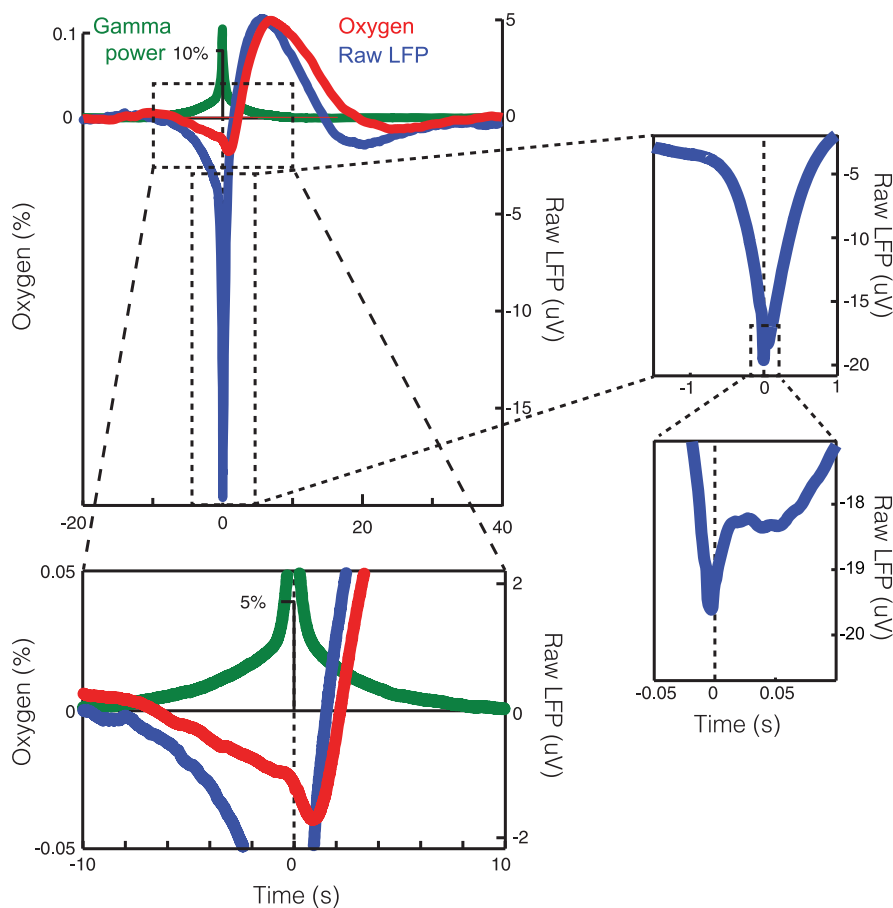


Fig. 2. Spike-triggered average of oxygen, LFP voltage and gamma-band LFP. Spikes are aligned at time zero. The top left shows a 60 s (60,000 ms) window around the spike, with expanded views in the insets. Raw LFP (blue) and oxygen (red) show a slow drop starting ~ 8 s before the spike. LFP has a strong negative transient around the time of the spike that lasts for ~ 2 s. Riding on top of this is a small (~ 1 uV), brief (< 20 ms) negativity almost exactly coincident with the spike (insets on right). This brief negativity likely represents contamination of the high frequency spike onto the LFP signal, while the bulk of the negativity has a duration that is too long (more than 1000 times the duration of a spike) to be explained in this way. After the spike, oxygen and raw LFP each show a delayed increase that peaks around 6 s, slowly falls back toward the baseline, and then overshoots and becomes slightly negative for ~ 10 s. The late slow responses in oxygen and raw LFP closely match each other, with raw LFP leading by about 0.5 s. Gamma LFP power (green) rises slowly prior to the spike, peaks abruptly at the time zero, then shows a symmetric decrease (inset on bottom left). Power is elevated from ± 8 s around time zero, though the strongest response occurs within ± 1 s of the spike (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

range of electrophysiology signals and oxygen. After low-pass filtering most of these signals, including in particular unit activity and gamma power, they still lead oxygen by at least 6 s. In contrast, applying the exact same low-pass filter to the LFP voltage signal yields a signal which leads oxygen by only ~ 0.5 s. Thus, the temporal coupling between electrophysiological signals and oxygen is tightest for infraslow LFP.

To further explore the temporal relationships between spikes, LFP and oxygen, we computed the spike-triggered average of oxygen, as well as gamma and infraslow LFP (Fig. 2). Results were similar across regions (Fig. S7) and therefore combined. On average, oxygen increases rapidly one second after a spike, peaking after 6.1 s, falling back to baseline after 20 s and then undershooting baseline slightly for ~ 10 s. This bears some resemblance to the transfer functions that relate task-evoked gamma to oxygen concentration (Bentley et al., 2016) and to the BOLD signal itself (Logothetis et al., 2001). It also resembles the cross-correlation functions between (1) spontaneous BOLD and gamma BLP and (2) spontaneous BOLD and MUA reported by Shmuel and Leopold (2008), although there is controversy over whether their recordings captured spontaneous or evoked activity (Logothetis et al., 2009; Shmuel and Leopold, 2008). In contrast to oxygen, the spike-triggered average of gamma power shows a symmetric elevation that is greatest at the time of the spike (± 1 s) and returns to baseline within 8 s. Surprisingly, the spike-triggered average of raw LFP has a slow response that is almost identical to the oxygen response, though it leads (precedes) oxygen by 0.5 s. Similar results were obtained using either a 0.01 or 0.07 Hz high-pass cut-off for the LFP recording (Supplemental Fig. S5). The fact that oxygen slightly lags but otherwise tracks the infraslow changes in LFP, without the need for an intervening transfer function, is consistent with (but does not prove) a common cause or direct causal relationship between processes underlying the two signals. An alternative explanation for the similar response shapes, that low-frequency changes in field potential somehow affect

our oxygen measurement, can be ruled out by the fact that contamination would affect the oxygen and LFP measures nearly simultaneously rather than with the observed 0.5 s offset between the two.

The data are consistent with three scenarios: gamma LFP may mediate the spike-oxygen relationship, spikes may mediate the relationship between gamma LFP and oxygen (though considerable evidence argues against this possibility; see Logothetis, 2008 for discussion), or spikes and gamma LFP may each be independently related to oxygen through similar (and perhaps overlapping) mechanisms.

On a separate note, starting 8 s prior to the spike, both raw LFP and oxygen drop below their baseline levels (Fig. 2, inset). This could indicate that the processes underlying one (or both) of these signals influences spike probability, however the pre-spike drop in oxygen is not consistent across regions (Fig. S7). Raw LFP shows an abrupt depolarization that begins one second before the spike and lasts for about two seconds. The duration of this depolarization is far too long to reflect the electrical activity of the spike itself (Fig. 2, insets to right). It may instead reflect the tendency of spikes to occur in so-called “up-states”, that is, periods in which the intracellular potential is elevated close to threshold and the extracellular potential is depressed (Wilson, 2008).

Immediately after the spike there is a small sudden drop in oxygen (lower left inset) that is superimposed on the slower drop. This could reflect a local and transient drop in oxygen similar to that reported by Thompson et al. (Thompson et al., 2003).

2.2. Spikes predict infraslow LFP, which in turn predicts oxygen

Lagged correlations (Fig. 1C) and spike-triggered averages (Fig. 2) suggest the relationship between infraslow LFP and oxygen is lower latency and less complex than that between spikes and gamma LFP and oxygen. This observation is consistent with several causal explanations,

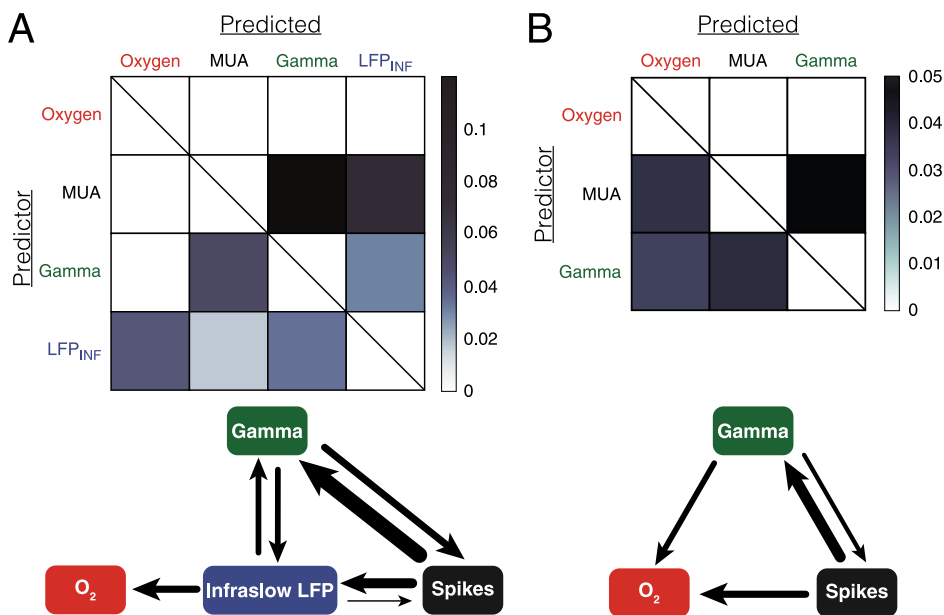


Fig. 3. Multivariate Granger causality analysis between oxygen and electrophysiological signals. Each cell shows whether forecasts of its predicted signal (column) improve when its predictor signal (row) is included in the model alongside the rest of the signals (see Methods). (A) Top, oxygen does not improve forecasts of any electrophysiological signal (first row), only infraslow LFP (LFP_{INF}) improves oxygen forecasts (first column), infraslow LFP forecasts are most improved by MUA (fourth column), and MUA improves forecasts of every other signal except oxygen (second row). Bottom, a graph of the Granger causality result. Arrow widths correspond to the magnitude forecast improvement, which is quantized into five categories (> 0.1 , $0.1-0.05$, $0.05-0.04$, $0.04-0.03$, $0.03-0.02$; for Granger causality > 0.02 , $p < 0.01$). Weak improvements (< 0.02) are not shown. (B) When infraslow LFP is excluded from the analysis, MUA and gamma both improve oxygen forecasts (first column). Data are from V3 and PCC; see Supplemental Fig. 8 for region-specific responses.

none of which can be definitively ruled out without (future) interventional experiments. To better understand which explanations are more likely based only on the current data, we examined the unique predictive utility of each signal with respect to all of the other signals using multivariate Granger causality analysis (Fig. 3, see Methods). MUA improves forecasts of LFP and gamma power. Oxygen forecasts are most improved by infraslow LFP ($p < 0.001$, with correction for multiple comparisons), while infraslow LFP forecasts are most improved by spikes ($p < 0.001$, corrected for multiple comparisons). It is of note that a connection between gamma power and oxygen appears only when infraslow LFP is excluded from the analysis (Fig. 3B). Multivariate Granger discounts predictability that is not unique to a particular signal. Thus if both infraslow LFP and gamma power predict future oxygen, then to the extent that some of that predictive power is common to the two signals, the predictive power will not be reported (see Methods). Finally, infraslow LFP has a small feedback effect on spikes and gamma LFP, which could correspond to the infraslow LFP change that precedes the spike in the spike-triggered average. The key findings were similar across regions and robust to 10-fold changes in the temporal window of the analysis (Supp. Fig. 8).

2.3. Oxygen correlation tracks correlation in the spiking activity of neurons

Next we asked whether electrophysiological signals, especially single units and MUA, show long-range, network-aligned correlations similar to those found previously in BOLD and in polarographic oxygen (within-modality correlations). We distinguish two types of long-range correlations. Within-network correlations are measured within a single network, e.g., between left and right PCC, or between left and right V3. Across-network correlations are measured between two different networks, e.g. between left PCC and right V3. We found, in agreement with a previous report (Li et al., 2015), that within-network correlations are mostly similar to each other (oxygen correlation[*left PCC-right PCC*]= 0.51 ± 0.08 [Pearson's r], oxygen correlation[*left V3-right V3*]= 0.41 ± 0.07 ; oxygen correlation[*left IPS-right IPS*]= 0.29 ± 0.08 ; p (PCC vs. V3)= 0.2 , p (IPS vs. V3)= 0.1), and that the two across-network correlations, within and across hemispheres, are also similar to each other (oxygen correlation[V3-PCC, *within hemisphere*]= 0.24 ± 0.04 , oxygen correlation[V3-PCC, *across hemisphere*]= 0.21 ± 0.04 ; $p = 0.2$). We therefore group our results into “within-network” and “across-network” (results by region are shown in Supplemental Fig. S9). We considered

a range of electrophysiological signals, including raw LFP, band-limited LFP power, multiple unit and single unit activity. We do not apply global signal regression (see Methods), which is commonly used in rsfMRI analysis to reduce widely shared variance arising from a mixture of neuronal, non-neuronal and artifactual effects (Murphy et al., 2009; Murphy and Fox, 2017; Power et al., 2014, 2018; Schölvinck et al., 2010). It remains unclear how the relationship between fMRI data and neural activity is affected by global signal regression.

Fig. 4A shows that raw LFP shows significant correlation across regions, and that within-network correlation (red) is significantly stronger than cross-network correlation (blue). Correlations are present and significantly greater than zero from 0.006 to 200 Hz. Correlation peaks at 0.05 Hz and rolls off gradually at higher frequencies. Correlation drops steeply for frequencies below 0.01 Hz, which may be due in part to high-pass filtering with a -3 dB point at 0.07 or 0.01 Hz (see Materials and Methods, Recording). Within-network correlation is substantially and significantly greater than cross-network correlation from 0.03 to 8 Hz, with small but significant differences persisting up to 30 Hz. One can also consider correlations in LFP power. In computing LFP power, there are two different frequencies that need to be taken into account. First, there is the underlying frequency at which we measure power. A pure 60 Hz sine wave, for example, has power at 60 Hz but not at any other frequency. Second, we can consider changes in power at a particular temporal scale. The temporal scale we are most interested in is that of long-range oxygen correlations (MRI-based resting-state functional connectivity), that is, the infraslow range from 0.01 to 0.1 Hz. We therefore consider band-limited power, computed over a small range of frequencies, and then filtered to consider only changes in the power envelope that occur in the infraslow range. Fig. 4B shows that band-limited LFP power, like oxygen and like raw LFP, shows significant long-range correlation, and that this correlation is larger within-network (red) than across-network (blue). The correlation difference is substantial and statistically significant for frequencies from 0.04 to 200 Hz. (The maximum frequency was limited by the LFP sampling rate.) The correlation increases with frequency up to around 55 Hz (within the gamma band) and falls off for still higher frequencies. To summarize, correlation in the raw LFP is strongest at low frequencies, centered around 0.05 Hz, while correlation in band-limited LFP power is strongest at high frequencies, centered around 55 Hz. Such correlations within raw LFP and band-limited LFP are robust across our recording configurations (Supplemental Fig. S3). Several studies have looked at non-linear

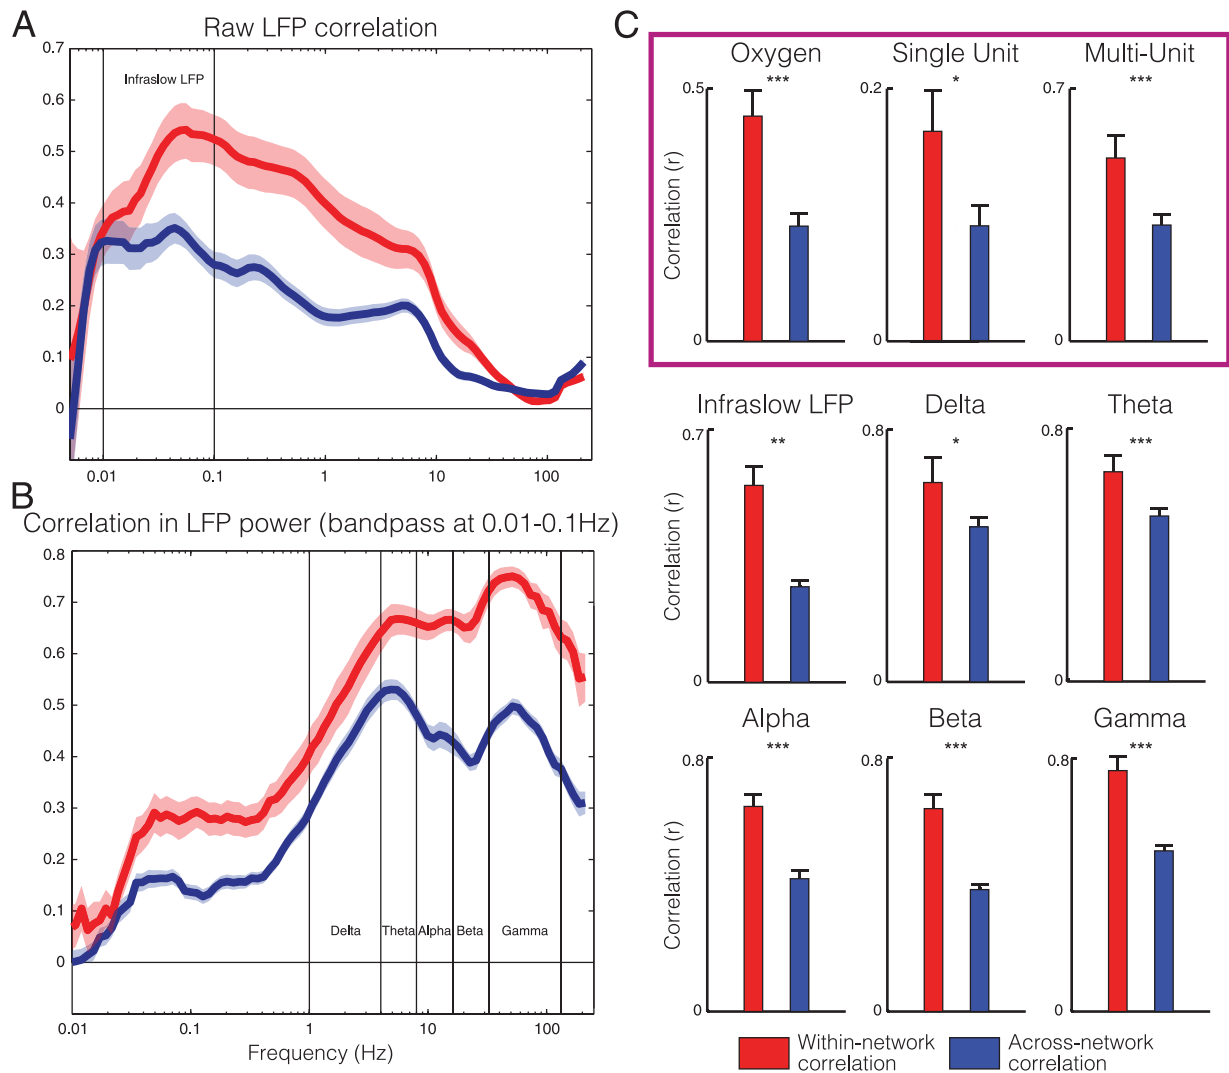


Fig. 4. Within-mode electrophysiological and oxygen correlations. (A). Long-range correlation in raw LFP signal by frequency, within-network (red) versus across-network (blue). Correlation is present in a wide range of frequencies (0.006–200 Hz). Peak correlation is around 0.05 Hz (within-network: 0.055 Hz, across-network: 0.049 Hz). (B). Long-range correlation in LFP power. LFP power was computed at half octave bands and the power at each band was then filtered to 0.01–0.1 Hz and used to compute long-range correlations. The strongest correlation was in the gamma band (55 Hz). Within-network correlations were greater than across-network correlations at all frequencies. (C). Like LFP, single and multi-unit activities show significant long-range, infraslow network-dependent correlations (top center and right) similar to those in oxygen (top left) and LFP (lower two rows). Correlations are significantly higher for within-network than across-network correlations in every case (oxygen: $t(151) = 4.2365$; single-unit: $t(78) = 1.997$; multi-unit: $t(151) = 7.1176$; infraslow LFP: $t(151) = 3.2154$; delta: $t(151) = 2.2663$; theta: $t(151) = 3.6047$; alpha: $t(151) = 5.5918$; beta: $t(151) = 8.2037$; gamma: $t(151) = 9.2618$; for single-units $n = 20$ within-network and $n = 60$ across-network; for all other signals $n = 34$ within-network and $n = 119$ across-network) and are significantly lower for single units than for multi-units (note different scales). For all 9 signals shown in C, correlations are computed within the frequency range of 0.01–0.1 Hz. *** = $p < 0.001$; ** = $p < 0.005$; * = $p < 0.05$. Data are combined across regions (see Fig. S9) (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

relationships between LFP bands, as well as cross-frequency coupling, the relationship between the low frequency raw LFP signal and the high frequency LFP band-limited power (He et al., 2010; Ko et al., 2011). Though interesting, these topics are beyond the scope of the current study.

Unit correlations were computed using firing rates filtered to 0.01–0.1 Hz. Fig. 4C shows that long-range infraslow correlations are present in both MUA and single units, and that the unit correlations show a spatial pattern similar to long-range low-frequency correlations in oxygen polarography (top row), as well as raw LFP and band-limited LFP power (middle and lower rows). In particular, like oxygen correlations, both within- and across- network unit correlations were positive ($p_{\text{within}} < 0.001$ and $p_{\text{across}} < 0.001$) and consistently higher within-network than across-network (oxygen, $p < 0.001$; single units, $p < 0.05$;

MUA, $p < 0.001$) (Li et al., 2015; Murphy et al., 2009). The reduced correlation of single units compared to multiple units ($r = 0.16$ versus 0.55 for within-network) is to be expected, since the MUA signal is pooled over many individual single units. The mean single unit correlation would be comparable to the multiunit correlation only if every neighboring single unit carried identical or nearly identical signals; this is clearly not the case.

Next we asked whether correlations in unit activity occur within the same frequency range as correlations in oxygen. Fig. 5A shows that, like long-range oxygen correlations, both single-unit and multi-unit correlations are band-limited. This is consistent with our previous results, though the literature is mixed on this point (see Discussion in Li et al. 2015). The peak correlations are similar, though slightly lower in frequency for units (0.03 Hz for MUA, 0.04 Hz for single units) com-

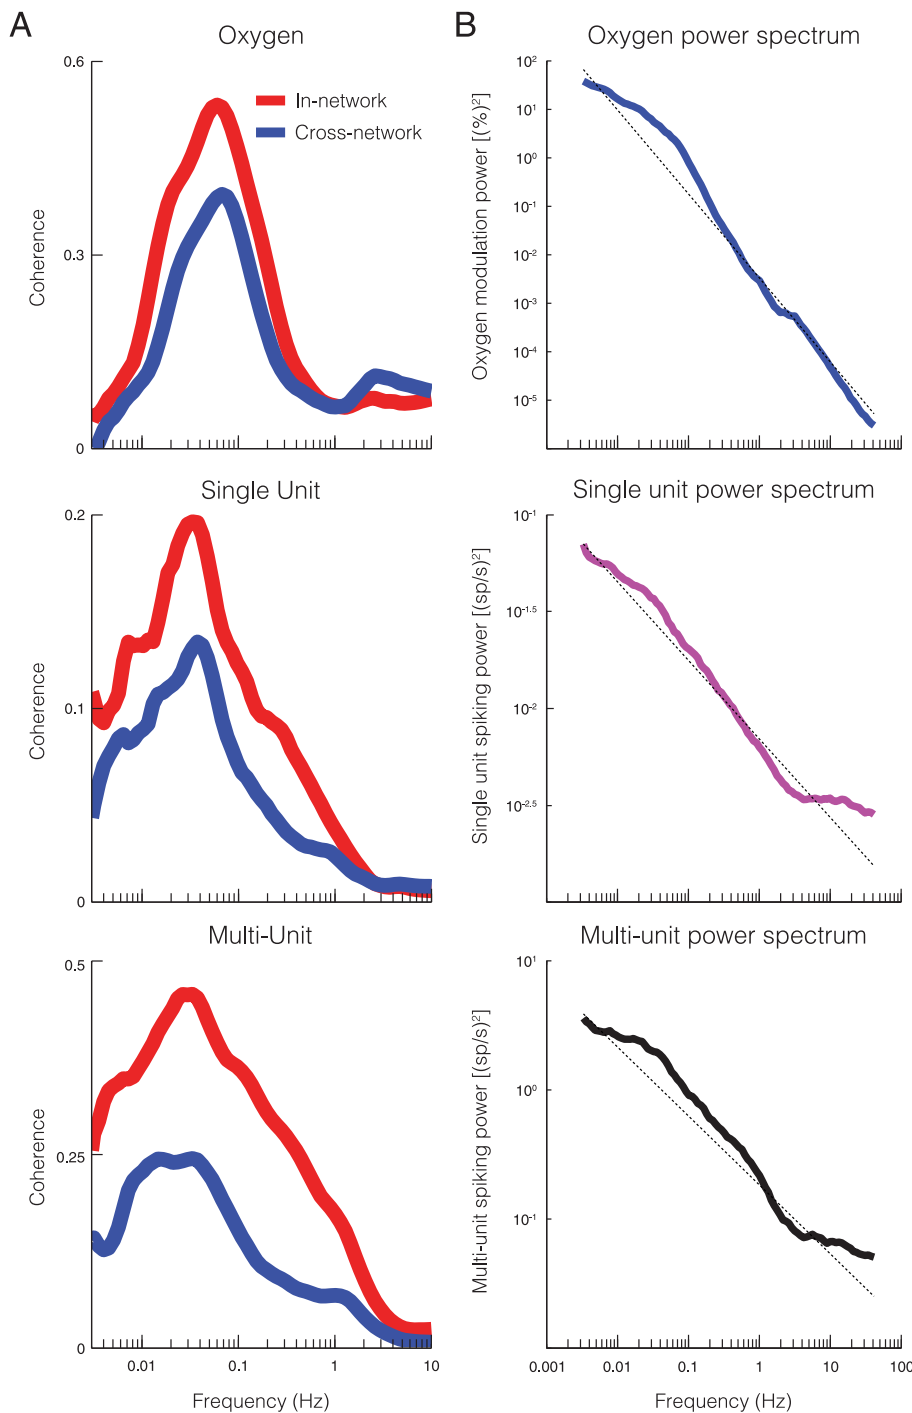


Fig. 5. Frequency distribution of correlation and power in units and oxygen. (A). Oxygen and unit coherence as a function of frequency. Coherence provides a measure of the frequency distribution of the correlation. Both within- and across-network coherences are band-limited. This is true for oxygen, single unit, and multi-unit coherences. (Unlike LFP there was no high-pass filtering of spikes nor oxygen signals that might explain the fall-off below 0.01 Hz.) The band-limited coherences suggest that a rhythmic or near-rhythmic mechanism drives long-range, network-aligned correlation. Unlike either unit coherence or gamma BLP (Supplemental Fig. 11), whose long-range coherences are significant up to 2 Hz, oxygen coherence does not extend above 0.3 Hz. This difference is consistent with the idea that there is a low-pass stage in between neuronal activity (spikes) and the hemodynamic response. Interestingly, merely low-passing the neuronal signals, including in particular single unit activity and gamma power, still leaves us with a ~6 s lag (Fig. 1C). In contrast, using the exact same low-pass filter on the raw LFP signal leads to a single with ~0 s lag.. (B). Power spectra of oxygen, multi-unit activity, and single unit activity. In all three cases, the power increases as frequency drops. The relationship between power and frequency is well fit by a linear relationship between the log of each signal [$\log(\text{power}) = -\beta \log(\text{frequency}) + k$]. This is typically referred to a 1/f power spectrum. There are deviations from 1/f, most pronounced at around 0.01 to 0.1 Hz, which could reflect the process(es) driving the band-limited correlation. Above 10 Hz there is a plateau in spike power, which may reflect a lower limit on spike power for our recording conditions. Power above 10 Hz is excluded from the estimation of the 1/f relationship.

pared to oxygen (0.06 Hz). Unit correlations extend farther in both directions compared to oxygen, for a full-width at half-max bandwidth of about 1.5 decades for oxygen, 2 decades for single spikes and 3 decades for MUA. The attenuation of oxygen correlation at the high end of this band, relative to spike correlation, is consistent with the high-cut filtering characteristics of the hemodynamic transfer function (Bentley et al., 2016; Logothetis et al., 2001). The attenuation of oxygen correlation at the low end of this band, and the shift to a higher peak correlation frequency compared to unit activity, are consistent with a low-cut filtering effect, perhaps related to the post stimulus undershoot in the hemodynamic response function. However, apart from this, the patterns of single-unit, multi-unit and oxygen correlations are all roughly similar,

with the highest correlations occurring in the range from about 0.01 to 0.1 Hz.

The similarity between the frequency distributions of spike and oxygen correlations might easily be explained if spike and oxygen signals were themselves similarly band-limited. This is not the case, however. For spikes as well as oxygen, the local power drops with increasing frequency over a very wide range (Fig. 5B). Therefore the similarities between long-distance spike correlation and long-distance oxygen correlation are not inherited from the local power spectra. Furthermore, note that the peak in long-distance correlation in all three signals roughly corresponds to a point of maximum positive deviation from a $1/f\beta$ fit to the local power. This is consistent with a model in which there are

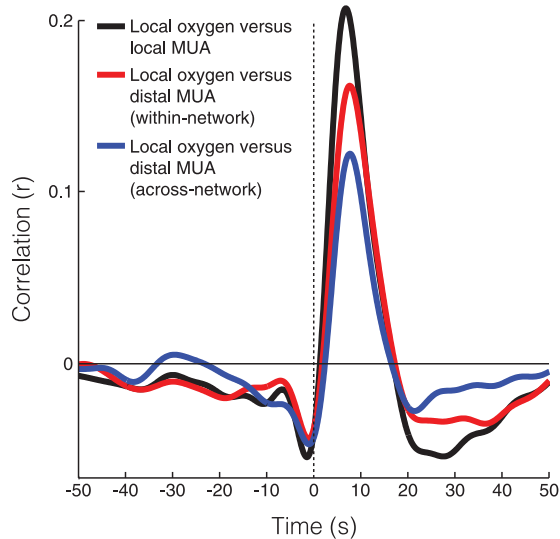


Fig. 6. Lagged correlation between oxygen and multi-unit activity within a site (local) and across sites separated by at least 1 cm (distal). Black: the mean lagged correlation between oxygen and multi-unit activity recorded from the same (local) site; red: the correlation between oxygen and multi-unit activity recorded from distal sites, both in the same network; blue: the correlation between oxygen and multi-unit activity recorded from distal sites, each from different networks. Data are combined across regions; for individual regions see Fig. S12 (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

two separate and at least partially independent sources of variance: one source or set of sources with only local influence and 1/f characteristics, and a second source or set of sources that influences signals at distal locations and has band-limited characteristics (Li et al., 2015).

Returning to the issue of whether oxygen correlations reflect correlations in neuronal activity, we directly compared oxygen and unit activity. To account for any hemodynamic delay, we computed lagged correlations. Local oxygen and local multi-unit activity are correlated with each other (Figs. 1 and 6, black), reaching a peak correlation at a lag of 6 s (spikes leading, oxygen following). Importantly, we also found that oxygen at one site is significantly correlated with multi-unit activity at distal sites with a similar delay (Fig. 6). The fact that spike-associated oxygen fluctuations occur simultaneously across sites support the view that correlated oxygen fluctuations arise in part due to correlated spiking activity. These cross-modal long-distance correlations are present both within and across networks. Local correlations are significantly stronger than distal correlation, and within-network distal correlations are stronger than across-network distal correlations (local vs. within network: $p < 0.05$, $t(232) = 2.6248$; within network vs. across network: $p < 0.05$, $t(302) = 7.7811$; statistics computed at a lag of 6 s). Together, these results strongly support the idea that oxygen correlation reflects correlated neural activity.

2.4. Infralow LFP accounts for long-range oxygen correlation

We next asked which electrophysiological signal – spikes, infralow LFP, or band-limited LFP power – best explains long-range oxygen correlations. Fig. 5 suggests that only a small, band-limited portion of the total local variance in oxygen (right column) is shared across distal sites (left column). The electrophysiological signals that best explain this restricted fraction of oxygen variance may differ from those that best explain the much larger local variations.

To address this possibility we developed a regression-based correlation dependence analysis. This analysis determines whether the shared variance between oxygen signals is the same as the shared variance be-

tween oxygen and electrophysiological signals by calculating how much oxygen correlations decrease after regressing out electrophysiological signals simultaneously recorded from the same locations (see Fig. 7 and Methods). This analysis indicates that the component of local oxygen variance underlying long-distance oxygen correlations is best explained by local infralow LFP, which accounts for 72% of within-network oxygen covariance and 57% of across-network covariance (Fig. 7). Gamma power is the next closest contender, but explains significantly less oxygen covariance than does infralow LFP ($P < 0.05$, paired t -test, with correction for multiple comparisons). (Note that error bars in the plot capture pooled variance, while statistics are computed pairwise and therefore are more powerful.) Multi-unit activity has the third strongest relationship to long-distance oxygen correlations.

3. Discussion

We measured correlates of infralow cortical tissue oxygenation in an awake monkeys at rest. Infralow LFP stands apart from all other electrophysiological signals as a correlate of oxygen, both at the local level (within an MRI voxel) and when comparing activity simultaneously recorded from distal sites within and across resting-state networks.

The local relationship between infralow LFP and oxygen is unique in several respects. Infralow LFP leads oxygen by ~ 500 ms, whereas spikes and LFP power (delta through gamma bands) lead oxygen by more than six seconds (Fig. 1). The short lead time of infralow LFP does not reflect cross-talk between our recording systems, which would be instantaneous. After accounting for time delays, infralow LFP is significantly more correlated with oxygen than any electrophysiological signal save gamma LFP power, which shows comparable correlation after a considerably longer delay (Fig. 1). Finally, the spike-triggered average (STA) of raw LFP features an infralow post-spike response that is nearly identical to the STA of oxygen, but again leads oxygen by ~ 500 ms (Fig. 2). Because these data were collected in the resting state rather than during a task, it is unclear to what extent the spikes that we record are clustered in time and space (regardless, we of course do not mean to imply that a single action potential determines the functional connectivity of the entire brain). Together, these findings strongly suggest a unique relationship between the processes and mechanisms underlying infralow oxygen and infralow LFP, as compared to other electrophysiological measures. Further, the lag between infralow LFP and oxygen is the same whether we consider spike-associated activity or full time-courses, consistent with overlap between the mechanisms relating infralow LFP and oxygen to each other and to spikes.

While we use an STA to reveal the post-spike responses, we do not mean to imply that single action potentials determine the functional connectivity of the entire brain. While these data were collected in the resting state rather than during a task, the single unit panel of Fig. 4C demonstrates that there is clustering of spikes in time and space. The relationship between infralow LFP correlations and long-range oxygen correlations also stands out. Infralow LFP accounts for most of the correlation in oxygen, significantly more than any other signal considered in this study. While we observed network-aligned long-range correlations in all electrophysiological signals, correlation is highest at infralow LFP frequencies and in infralow fluctuations of higher-frequency LFP power. Long-range spike correlations are also highest at infralow frequencies. Intriguingly, Granger causality analysis suggests that infralow fluctuations in multi-unit spiking activity drive colocalized infralow fluctuations in LFP, which in turn drive infralow fluctuations in oxygen. Together, these findings suggest that the neural activity and neurovascular coupling mechanisms relating oxygen and infralow LFP to each other and to spikes at the local level may underlie the bulk of long-range oxygen correlations.

Our findings stand in contrast to (but do not contradict) previous reports that identify gamma LFP (Goense and Logothetis, 2008; Logothetis et al., 2001; Shmuel and Leopold, 2008) and MUA (Shmuel and Leopold, 2008) as the best electrophysiological corre-

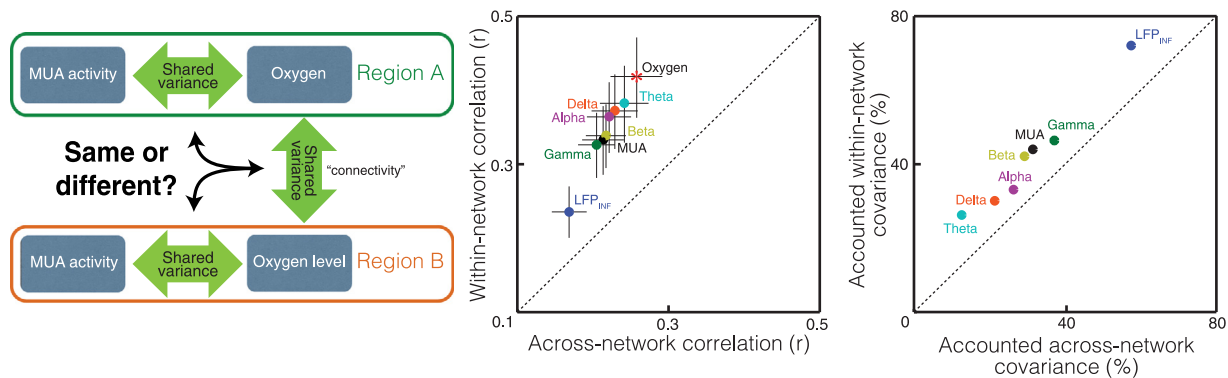


Fig. 7. Regression-based dependence analysis on oxygen correlation. *Left:* dependence analysis determines whether the shared variance between oxygen signals is the same as the shared variance between oxygen and MUA activity (or any other electrophysiological signal). *Middle:* oxygen correlation before (red asterisk) and after (colored circles) regressing out linearly transformed electrophysiological signals (see Methods). Regressing out infraslow LFP (LFP_{INF}) had the greatest effect on oxygen correlation, substantially reducing both within- and across-network correlation (within-network: $n = 27$; across-network: $n = 84$). *Right:* Infraslow LFP accounts for 72% of the covariance underlying within-network oxygen correlation, and 57% of the covariance underlying across-network correlation. Gamma power accounts for 46% of within-network oxygen covariance and 37% of across-network covariance. MUA accounts for 43% of within-network oxygen covariance and 31% of across-network covariance. Other electrophysiological signals account for only 10–42% of oxygen covariance. Data are combined across regions, for individual regions see Fig. S14.

lates of the BOLD signal. Gamma LFP power exhibits higher long-range within-mode correlation than infraslow LFP (Fig. 4c), and infraslow LFP and gamma LFP power are equally well correlated with local oxygen variance (Fig. 1c;), yet infraslow LFP best accounts for the relatively small fraction of local oxygen variance that underlies long-distance oxygen correlations (Fig. 7). This apparent dissociation between long-range oxygen and electrophysiological correlations could be explained by a neuronal process that drives long-range correlation of band-limited gamma power in the 0.01 to 0.1 Hz range, but has a less effect on oxygen levels than the processes underlying infraslow LFP. This is in turn consistent with the idea that different frequency components of electrical signals may reflect the activity of different neural processing pathways (Buice and Chow, 2013; Magri et al., 2012; Tong et al., 2015). Another relevant point is that gamma band power and infraslow LFP amplitude may be differentially contaminated by non-neural processes or by noise, which may then differentially affect both long-range correlations and correlations with local oxygen. In summary, although infraslow LFP does not exhibit the strongest within-mode long-range correlation among the electrophysiological signals, this does not preclude it from being the best predictor of long-range oxygen correlations.

There are many potential sources of infraslow LFP fluctuations, and it is likely that, at any given location, multiple such processes unfold more or less independently of one another. Here, we focus only on those that offer a clear potential explanation of the main observations described in this report. The infraslow LFP could reflect an elevation in astrocyte calcium waves, which have been proposed to help control blood flow (Hillman, 2014). This could explain the strong coupling and relatively short delay between infraslow LFP and oxygen changes. This explanation requires that astrocytic processes be spatially aligned, else the voltage gradients would cancel one another and not be detected by our electrodes. Alternatively, infraslow LFP could reflect prolonged depolarization of pyramidal apical dendrites, which are neatly aligned to form an open field arrangement capable of generating strong changes in potential (Mitzdorf, 1985). If true, this would imply that infraslow LFP signals reflect inputs that drive changes in excitability, which then manifest as changes in spiking activity. However, in the resting state, spike-triggered averaging and Granger causality analysis show that spiking activity leads and may drive the infraslow LFP; the reverse (slow LFP influencing spike rate) may occur but is at best only a weak effect (Schmidt et al., 2014).

Our data are most consistent with infraslow LFP reflecting a process that occurs following spikes. For example, infraslow LFP could corre-

spond to a membrane hyperpolarization that arises from a slow recovery of voltage-gated sodium channels from the inactivation state following action potential generation (Mickus et al., 1999; Ogata and Tatebayashi, 1992; Toib et al., 1998). However, this interpretation does not offer a clear explanation of the close correspondence between infraslow LFP and infraslow oxygen. Future studies combining intracellular and extracellular recordings with pharmacological approaches to manipulate channels may elucidate the physiological processes underlying both signals, as well as their relationship to brain function.

Our principle findings have clear antecedents in the literature. We were able to show that spikes, like oxygen, show long-range correlation (Fig. 4) because we considered correlation on a time scale of seconds. Many previous studies focused instead on a scale of milliseconds, and as a result are blind to slow correlations. However, long-range slow spike-spike correlations have been previously reported in human bilateral auditory cortices (Nir et al., 2008). Ma et al. (2016) described a close relationship between correlated bilateral fluctuations of excitatory neuron calcium levels and optically recorded hemoglobin (Ma et al., 2016). Similarly, several studies have anticipated our finding that the infraslow LFP signal is well correlated with oxygen (He et al., 2008; Li et al., 2014). This relationship has eluded wider observation in part because of a long tradition of looking at EEG power rather than the raw (unfiltered) LFP signal, and in part because most off-the-shelf systems for recording LFP incorporate high-pass filters that eliminate infraslow LFP. Finally, we were able to more directly establish the relationship between infraslow LFP and correlated oxygen (Fig. 7) because we focused our analyses on the correlated oxygen signal. The correlated oxygen signal differs substantially from the local oxygen signal; the former has bandpass characteristics while the latter is $1/f$ (Fig. 5, top row). While gamma power can be linearly transformed (via a hemodynamic function) to resemble the local oxygen signal, it has a significantly weaker relationship than infraslow LFP with long-range oxygen correlation. In contrast, infraslow LFP has a strong relationship with both local oxygen levels (Figs. 1–3) and long-range oxygen correlation (Fig. 7).

4. Materials & methods

4.1. Animals and behavior

Three macaques served as subjects in this study. Animals were cared for and handled in accordance with the Guide for the Care and Use of Laboratory Animals, and all procedures were approved by the Washington University Animal Studies Committee. During recording, macaques

were fully hydrated and sat head-fixed in a dark room. Behavior was unconstrained, and the animals had no expectation of a task or reward. Animals naturally relaxed in the setup. Eyes were partially or fully closed $14.7 \pm 6.4\%$, $28.4 \pm 7.3\%$ and $85.5 \pm 2.1\%$ of the time, respectively, for monkeys P, E and L.

4.2. Recording

A total of 65 sessions were recorded with an average duration of 50 min. Intracranial oxygen and electrophysiological signals were recorded simultaneously from both hemispheres (**Supplemental Fig. S15**). Some sessions (19 from monkey P, 16 from monkey L) recorded bilateral PCC and bilateral V3 sites (4 sites total). The remaining 30 sessions (from monkey E) recorded bilateral PCC and bilateral IPS sites (4 sites total). At each site, two electrodes (separated by 0.5 mm) recorded concurrent oxygen and electrophysiological signals separately. Results were generally similar for the three monkeys, and thus the data were combined.

Electrodes were targeted to each area of interest using anatomical MRI images and physiology (Bentley et al., 2016). Briefly, each animal's brain was accessed via bilateral 15 mm (internal diameter) chronic custom recording chambers. T1 weighted MRI images (MPRAGE; 0.5 mm isotropic voxels) were obtained using a custom phantom in the chamber that provides visualization of the chamber and allows for the virtual projection of a chamber-based coordinate system down into the brain. In one monkey, two small manganese injections were used to validate targeting accuracy by comparing their expected locations to their actual locations in an MRI image acquired directly following the injections. Prior to data collection, boundaries for PCC, V3 and IPS recording regions were defined on the MRI image. The positions of PCC, V3 and IPS were further validated based on their respective oxygen and electrophysiological responses to visual stimulation, which were recorded immediately before the resting state data (Bentley et al., 2016).

The multi-channel voltage clamp system used to record oxygen signals is described and validated in Li et al. (2015) and Bentley et al. (2016), and a similar system is described in Thompson et al. (2003) and (Khine et al., 2003). Oxygen signals (the clamp's current) were sampled without filtering and stored at 1 KHz, then mean-centered offline before analysis. Electrophysiological signals were recorded using tungsten microelectrodes (Alpha Omega LTD). Li et al. (2015) includes example traces of both signals (Li et al., 2015). Electrical activity was filtered, amplified and sampled using the Plexon MAP system (Plexon, Inc.). For identifying action potentials, data were filtered from 400 Hz (high pass, 2-pole analog filter) to 8000 Hz (low pass, 6-pole filter) and stored at 40 KHz. For local field potentials (LFP), the data presented in the main text and figures (54 recording sessions) used a 200 Hz 2-pole low-pass filter and 0.07 Hz 1-pole high-pass analog filter, and were digitized at 1 KHz. A high-pass filter is necessary because LFP power is roughly inversely related to frequency, and so very low frequencies will saturate the amplifiers. A 1-pole filter reduces the signal amplitude by half when the frequency halves, and so prevents saturation while still passing through a large amount of low frequency signal. As a result, we can measure (attenuated) LFP amplitude down to at least 0.001 Hz in this system. However, a low pass filter will impose a phase shift, with larger impact for lower frequencies. This could affect the raw LFP analyses associated with Figs. 1, 2 and 7 (LFP power, oxygen and single unit activity are not affected by this cut-off). To minimize this distortion, we replaced the 0.07 high pass filter Hz filter with a 0.01 Hz 1-pole analog filter, collected 11 additional data sets from monkey L, and re-ran these analyses. These results were essentially identical to the 0.07 Hz data, and are presented in Supplemental Results. A second reason to use a high-pass filter is to prevent slow drift in signals. Despite our unusually low high-pass cut-offs, we saw little evidence for slow drift, and no evidence of a significant drift in our recordings (Supplemental Fig. S1).

Oxygen and electrophysiological recording systems were electrically isolated from each other using a custom-built optical isolation system. This system limits possible cross-talk mechanisms to capacitive coupling to earth. Isolation was confirmed *in vitro*.

4.3. Analysis

All analyses were performed with custom software written in Matlab (MathWorks). T-tests were all two-tailed. Oxygen polarographic signals, like BOLD signals, reflect relative rather than absolute oxygen levels. Therefore the polarographic signals, like BOLD data, were expressed as percentage deviation from the mean signal level. High-frequency electrical activity was analyzed offline to extract identifiable action potentials (single unit activity) or multi-unit activity (MUA), applying clustering algorithms in a feature space based on principle component projections (Offline Sorter, Plexon Inc.). The low-frequency electrical activity ("raw LFP") was notch filtered at 60 Hz to remove power line noise. The raw LFP signal was used for Figs 1A, 2 and 4A. The raw LFP signal was further processed to extract band-limited LFP power. To extract band-limited power, the LFP was first digitally filtered into five bands: delta (1–4 Hz), theta (4–8 Hz), alpha (8–16 Hz), beta (16–32 Hz), and gamma (32–128 Hz) (We use these bands because they are established in the EEG literature, and not because we observed well-separated rhythms with distinct dynamics corresponding to these particular bands.) A Hilbert transform was applied to compute the power. Finally, the Hilbert-derived power was digitally band-pass filtered to 0.01 to 0.1 Hz. To extract infraslow LFP amplitude, the raw LFP signal was digitally band-pass filtered to 0.01–0.1 Hz.

4.4. Heart-rate removal

Cardiac and respiratory activity both drive variance in polarographic oxygen signals. We estimated and regressed out this variance session-by-session and site-by-site, following the same procedure described in (Li et al., 2015).

4.5. Filtering

Most studies report that MRI functional connectivity is most robust at slow time scales from about 0.01 to 0.1 Hz (Salvador et al., 2008; Sasai et al., 2011). We previously showed, and replicate in this study, that the long-distance correlation in the polarographic oxygen signal lies within this same frequency range (Li et al., 2015). Thus many of our analyses compare electrophysiological signals to infraslow (0.01 to 0.1 Hz) filtered oxygen signals. This raises a potential confound, explained below.

A number of our analyses focus on the timing of the relationships between electrophysiological signals and oxygen, essentially asking: after a change in signal A, how long is the delay before the corresponding change in oxygen? On the face of it, this is a straightforward question. But there is a hidden complication. In most of our analyses, signal A and oxygen have both been passed through at least one filter, and filtering a signal can shift it in time. If oxygen has been shifted forwards in time by a low-pass filter but signal A has not, then the delay between changes in signal A and related changes in oxygen will increase. This would make analysis of the timing of the relationship between the two signals hard to interpret (perhaps meaningless). The problem is actually worse, because we compare a number of filtered signals to oxygen with the goal of asking how their temporal relationships to oxygen differ, and what those differences suggest about the causal relationships driving changes in oxygen. But if the timing differences arise in large part due to the different filters applied to each signal, interpreting such results becomes quite difficult if not impossible.

To avoid this confound, we digitally filtered the electrophysiological signals (including unit activity, raw LFP, and LFP power) between

0.01 to 0.1 Hz using the same type-II Chebyshev filter applied to oxygen data before all analyses that use infraslow oxygen, including those illustrated by Figs 1B,C, 3, 4A–C, 6, and 7. For single and multi-unit activity, this filtering was applied to the time series representing firing rate. This filtering also eliminates any confound that might otherwise arise due to differences in the frequency content of the electrophysiological signals and infraslow oxygen. A similar matching of frequency content can be achieved by averaging the signal within sequential time bins (Magri et al., 2012). However, filtering is a more robust procedure in this regard than binning (Oweiss, 2010).

4.6. Correlation and coherence

Correlation was calculated as Pearson's r . Lagged correlation between electrophysiological signals and oxygen was obtained by shifting the electrophysiological signals forwards or backwards in time before calculating their correlation with oxygen. The methods used to calculate coherence and the mean expected bias in coherence are described in Li et al. (2015). For the result reported in Fig. 5, the mean expected bias was subtracted out so that a value of zero corresponds to the consistency of phase differences that would be obtained by chance were the null hypothesis true (no coherence).

4.7. Regression-based dependence analysis

In order to determine which signals drive long-range correlations, one might ask if the correlated portion of two putative driver signals matches the correlated portion of two putative driven signals. For example, given a gamma LFP signal and an oxygen signal from site A, and a gamma LFP and oxygen signal from site B, one might ask if the correlated portion of the two gamma LFP power signals predicts (matches) the correlated portion of the two oxygen signals. Unfortunately, although one can compute a correlation coefficient between two signals, one cannot accurately extract the correlated signal. The extraction of a correlated component assumes that each of the two individual signals can be separated into a shared (correlated) component and an unshared (non-correlated) component. Thus the two original individual signals are to be separated into three components – one correlated and two uncorrelated. Since there are more components than sources, there is no closed form solution to this problem. A similar argument can be made for three signals, four signals, etc.

As the number of signals becomes very large, a component that closely approximates the correlated component can be extracted using methods such as principle component analysis or independent component analysis. The general idea behind these methods is to construct a coefficient matrix and then use that matrix to determine a weighted average of the original signals, which constitutes a pseudo-correlated component (Li et al., 2009). Methods differ on the exact process used to construct the coefficient matrix. Regardless of the method, the extracted pseudo-correlated component is a weighted average of the original signals, and so the extracted component is inevitably contaminated to some degree by the uncorrelated components of the original signals. The amount of contamination depends in part on the magnitude of the correlation (which reflects the ratio between the correlated and non-correlated components in the original signals), and on the number of raw signals used to construct the correlated component.

Supplemental Fig. S2 shows that the amount of contamination decreases as the number of raw signals increases. In the case of our data, we have two sets of two within-network signals (each set corresponding to one of two functional networks) and four sets of two across-network signals. The extracted pseudo-correlated component is heavily contaminated by local un-correlated signals (38.6% for the within-network component, and 63.4% for the across-network component). To reduce contamination below 10%, at least 12 within-network signals and at least 33 across-network signals would be required (Supplemental Fig. S2). (NB

this means our data cannot provide an uncontaminated “global signal” for use in global signal regression).

With the difficulty of extracting the correlated component in mind, we developed a regression-based correlation dependence analysis to determine which electrophysiological signal best accounts for the correlation we see in oxygen signals. Correlation dependence analysis asks how much the correlation between one pair of signals (a putative driven signal) may depend on the correlation between another pair of signals (a putative driver signal). As a simplified example, consider recording oxygen, gamma and beta signals from sites A and B. Suppose that all three signals show within-mode correlation, and that the correlated component of each site's oxygen signal is linearly predicted by their respective gamma (but not beta) signals. In this case, regressing out A_{gamma} (gamma power recorded at site A) from A_{oxygen} and regressing out B_{gamma} from B_{oxygen} will eliminate the correlation between A_{oxygen} and B_{oxygen} . In contrast, regressing out A_{beta} from A_{oxygen} and B_{beta} from B_{oxygen} (respectively) will have no effect on the correlation between A_{oxygen} and B_{oxygen} . More generally, the extent to which interregional oxygen correlation diminishes after regressing out a particular electrophysiological signal reveals the extent to which the correlated component of oxygen signals statistically depends on correlated activity reflected by that particular electrophysiological signal.

To determine the electrophysiological dependencies of oxygen correlation, we first compute the baseline covariance of oxygen correlation (0.01–0.1 Hz) using Pearson's correlation ($\text{Covariance}_{\text{[Before regression]}}$). We then ask how much of this baseline covariance can be accounted for by electrophysiological signals. This is accomplished using linear regression after applying transfer functions to the electrophysiological signals, as described in Logothetis et al. (2001). A transfer function is computed separately for each electrophysiological signal (session-by-session and site-by-site) via the Welsh method (Oweiss, 2010). Briefly, the data are broken into 5-minute half-overlapping windows. Oxygen and an electrophysiological signal recorded from an adjacent electrode are converted into complex exponential form using Fourier transformation ($A_O(t,f)e^{i\omega O(t,f)}$ and $A_E(t,f)e^{i\omega E(t,f)}$, respectively). A transfer function representing the relationship between the electrophysiological signal and the oxygen level is computed as follows:

Transfer Function

$$= \frac{[\sum A_O(t,f)e^{i\omega O(t,f)} * A_E(t,f)e^{i\omega E(t,f)}] - [\sum A_O(t,f)e^{i\omega O(t,f)}] * [\sum A_E(t,f)e^{i\omega E(t,f)}]}{[\sum A_E(t,f)e^{i\omega E(t,f)}]^2 - [\sum A_E(t,f)e^{i\omega E(t,f)}]^2}$$

This transfer function is convolved with the electrophysiological signal and the result regressed out of the oxygen signal, thereby removing from the oxygen signal all variance that can be predicted by the electrophysiological signal. Session-by-session and site-by-site computation allows for different transformations at different sites, accounting for possible site-specific variability in neurohemodynamic relationships (Bentley et al., 2016).

The final step in correlation dependence analysis is to compute the covariance in oxygen that remains after the influence of one electrophysiological signal has been regressed out ($\text{Covariance}_{\text{[After regression]}}$). The percentage of the total covariance in the oxygen that is accounted for by that electrophysiological signal is then simply:

$$\frac{(\text{Covariance}_{\text{Before regression}} - \text{Covariance}_{\text{After regression}})}{\text{Covariance}_{\text{Before regression}}} \times 100\%$$

We performed correlation dependence analysis using oxygen and electrophysiological signals filtered at 0.01–0.1 Hz.

4.8. Granger causality

Multivariate Granger causality was computed to gain further insight into the temporal relationships between signals (Barnett and Seth, 2014). Granger causality determines whether the history of signal A and B together can improve the prediction of the future of signal B, compared to using the history of B alone. If it does, then A is Granger causal to B. Multivariate Granger causality is an extension of this logic, where A is Granger causal to B, only when the history of signal A and

all other signals together improve the prediction of the future of B, compared to using signals excluding signal A. It is important to recognize that although Granger causality can provide evidence consistent with causality, it does not provide proof of causality.

Before computing Granger causality, we filtered both oxygen and electrophysiological signals between 0.01 and 0.1 Hz. We computed Granger causality using a range of model lengths (up to 10 s) to account for possible lags in the causal relationship between electrophysiological and oxygen signals. The results were similar across model lengths and recording locations (**Supplemental Fig. S8**), and we present the result for a model length of 1 s. Similar to transfer function computation, Granger causality was also computed session-by-session and site-by-site. Significance was tested by permuting data across recording sessions in order to generate a null distribution, against which we compared the actual results.

Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Code availability

The code that analyzed the neural data of this study are available from the corresponding author upon reasonable request.

Declaration of Competing Interest

The authors declare no competing interests.

Credit authorship contribution statement

Jingfeng M. Li: Visualization, Investigation, Formal analysis, Data curation, Writing – original draft. **Benjamin T. Acland:** Writing – original draft. **Alexander S. Brenner:** Investigation. **William J. Bentley:** Visualization. **Lawrence H. Snyder:** Visualization, Writing – original draft.

Acknowledgments

This work was supported by the [National Institute of Mental Health](#) at the National Institutes of Health (R01MH102471); the National Institute of Neurological Disorders and Stroke (R21NS105090 and R34NS118618), the National Institute of General Medical Sciences (T32GM008151) to JML, and the McDonnell Center for Cellular and Molecular Neurobiology.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.neuroimage.2021.118728](https://doi.org/10.1016/j.neuroimage.2021.118728).

References

Babapoor-Farrokhan, S., Hutchison, R.M., Gati, J.S., Menon, R.S., Everling, S., 2013. Functional connectivity patterns of medial and lateral macaque frontal eye fields reveal distinct visuomotor networks. *J. Neurophysiol.* 109, 2560–2570. doi:10.1152/jn.01000.2012.

Barnett, L., Seth, A.K., 2014. The MVGC multivariate Granger causality toolbox: a new approach to Granger-causal inference. *J. Neurosci. Methods* 223, 50–68. doi:10.1016/j.neumeth.2013.10.018.

Bentley, W.J., Li, J.M., Snyder, A.Z., Raichle, M.E., Snyder, L.H., 2016. Oxygen level and LFP in task-positive and task-negative areas: bridging BOLD fMRI and electrophysiology. *Cereb. Cortex* 26, 346–357. doi:10.1093/cercor/bhu260.

Birn, R.M., 2012. The role of physiological noise in resting-state functional connectivity. *Neuroimage* 62, 864–870. doi:10.1016/j.neuroimage.2012.01.016.

Buice, M.A., Chow, C.C., 2013. Dynamic finite size effects in spiking neural networks. *PLoS Comput. Biol.* 9, e1002872. doi:10.1371/journal.pcbi.1002872.

Felleman, D.J., Burkhalter, A., VanEssen, D.C., 1997. Cortical connections of areas V3 and VP of macaque monkey extrastriate visual cortex. *J. Comp. Neurol.* 379, 21–47. doi:10.1002/(Sici)1096-9861(19970303)379:1<21::Aid-Cnc3>3.0.Co;2-K.

Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Van Essen, D.C., Raichle, M.E., 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc. Natl. Acad. Sci. U. S. A.* 102, 9673–9678. doi:10.1073/pnas.0504136102.

Goense, J.B., Logothetis, N.K., 2008. Neurophysiology of the BOLD fMRI signal in awake monkeys. *Curr. Biol.* 18, 631–640. doi:10.1016/j.cub.2008.03.054.

Gorges, M., Roselli, F., Müller, H.P., Ludolph, A.C., Rasche, V., Kassubek, J., 2017. Functional connectivity mapping in the animal model: principles and applications of resting-state fMRI. *Front. Neurol.* 8, 200. doi:10.3389/fneur.2017.00200.

He, B.J., Zempel, J.M., Snyder, A.Z., Raichle, M.E., 2010. The temporal structures and functional significance of scale-free brain activity. *Neuron* 66, 353–369. doi:10.1016/j.neuron.2010.04.020.

He, B.Y.J., Snyder, A.Z., Zempel, J.M., Smyth, M.D., Raichle, M.E., 2008. Electrophysiological correlates of the brain's intrinsic large-scale functional architecture. *Proc. Natl. Acad. Sci. U.S.A.* 105, 16039–16044. doi:10.1073/pnas.0807010105.

He, Y., Wang, M., Chen, X., Pohmann, R., Polimeni, J.R., Scheffler, K., Rosen, B.R., Kleinfeld, D., Yu, X., 2018. Ultra-slow single-vessel BOLD and CBV-fMRI spatiotemporal dynamics and their correlation with neuronal intracellular calcium signals. *Neuron* 97, 925–939. doi:10.1016/j.neuron.2018.01.025, e925.

Hillman, E.M., 2014. Coupling mechanism and significance of the BOLD signal: a status report. *Annu. Rev. Neurosci.* 37, 161–181. doi:10.1146/annurev-neuro-071013-014111.

Hutchison, R.M., Gallivan, J.P., Culham, J.C., Gati, J.S., Menon, R.S., Everling, S., 2012. Functional connectivity of the frontal eye fields in humans and macaque monkeys investigated with resting-state fMRI. *J. Neurophysiol.* 107, 2463–2474. doi:10.1152/jn.00891.2011.

Khader, P., Schicke, T., Roder, B., Rosler, F., 2008. On the relationship between slow cortical potentials and BOLD signal changes in humans. *Int. J. Psychophysiol.* 67, 252–261. doi:10.1016/j.ijpsycho.2007.05.018.

Ko, A.L., Darvas, F., Poliakov, A., Ojemann, J., Sorensen, L.B., 2011. Quasi-periodic fluctuations in default mode network electrophysiology. *J. Neurosci. Off. J. Soc. Neurosci.* 31, 11728–11732. doi:10.1523/JNEUROSCI.5730-10.2011.

Kozberg, M.G., Hillman, E.M., 2016. Neurovascular coupling develops alongside neural circuits in the postnatal brain. *Neurogenesis* 3, e1244439. doi:10.1080/23262133.2016.1244439, (Austin).

Li, B., Liu, R., Huang, Q., Lu, J., Luo, Q., Li, P., 2014. Coherent slow cortical potentials reveal a superior localization of resting-state functional connectivity using voltage-sensitive dye imaging. *Neuroimage* 91, 162–168. doi:10.1016/j.neuroimage.2014.01.004.

Li, J.M., Bentley, W.J., Snyder, L.H., 2015. Functional connectivity arises from a slow rhythmic mechanism. *Proc. Natl. Acad. Sci. U. S. A.* 112, E2527–E2535. doi:10.1073/pnas.1419837112.

Li, K., Guo, L., Nie, J., Li, G., Liu, T., 2009. Review of methods for functional brain connectivity detection using fMRI. *Computerized medical imaging and graphics: the official journal of the. Comput. Med. Imaging Soc.* 33, 131–139. doi:10.1016/j.compmedimag.2008.10.011.

Logothetis, N.K., Murayama, Y., Augath, M., Steffen, T., Werner, J., Oeltermann, A., 2009. How not to study spontaneous activity. *Neuroimage* 45, 1080–1089. doi:10.1016/j.neuroimage.2009.01.010.

Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412, 150–157. doi:10.1038/35084005.

Ma, Y., Shaik, M.A., Kozberg, M.G., Kim, S.H., Portes, J.P., Timerman, D., Hillman, E.M., 2016. Resting-state hemodynamics are spatiotemporally coupled to synchronized and symmetric neural activity in excitatory neurons. *Proc. Natl. Acad. Sci. U. S. A.* 113, E8463–E8471. doi:10.1073/pnas.1525369113.

Magri, C., Schridde, U., Murayama, Y., Panzeri, S., Logothetis, N.K., 2012. The amplitude and timing of the BOLD signal reflects the relationship between local field potential power at different frequencies. *J. Neurosci.* 32, 1395–1407. doi:10.1523/JNEUROSCI.3985-11.2012.

Mars, R.B., Jbabdi, S., Sallet, J., O'Reilly, J.X., Croxson, P.L., Olivier, E., Noonan, M.P., Bergmann, C., Mitchell, A.S., Baxter, M.G., et al., 2011. Diffusion-weighted imaging tractography-based parcellation of the human parietal cortex and comparison with human and macaque resting-state functional connectivity. *J. Neurosci.* 31, 4087–4100. doi:10.1523/JNEUROSCI.5102-10.2011.

Mickus, T., Jung, H., Spruston, N., 1999. Properties of slow, cumulative sodium channel inactivation in rat hippocampal CA1 pyramidal neurons. *Biophys. J.* 76, 846–860.

Mitzdorf, U., 1985. Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiol. Rev.* 65, 37–100.

Murayama, Y., Biessmann, F., Meinecke, F.C., Müller, K.R., Augath, M., Oeltermann, A., Logothetis, N.K., 2010. Relationship between neural and hemodynamic signals during spontaneous activity studied with temporal kernel CCA. *Magn. Reson. Imaging* 28, 1095–1103. doi:10.1016/j.mri.2009.12.016.

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, 893–905. doi:10.1016/j.neuroimage.2008.09.036.

Murphy, K., Fox, M.D., 2017. Towards a consensus regarding global signal regression for resting state functional connectivity MRI. *Neuroimage* 154, 169–173. doi:10.1016/j.neuroimage.2016.11.052.

Murphy, M.C., Chan, K.C., Kim, S.G., Vazquez, A.L., 2018. Macroscale variation in resting-state neuronal activity and connectivity assessed by simultaneous calcium imaging, hemodynamic imaging and electrophysiology. *Neuroimage* 169, 352–362. doi:10.1016/j.neuroimage.2017.12.070.

Nir, Y., Mukamel, R., Dinstein, I., Privman, E., Harel, M., Fisch, L., Gelbard-Sagiv, H., Kipervasser, S., Andelman, F., Neufeld, M.Y., et al., 2008. Interhemispheric correla-

- tions of slow spontaneous neuronal fluctuations revealed in human sensory cortex. *Nat. Neurosci.* 11, 1100–1108. doi:10.1038/nn.2177.
- Ogata, N., Tatebayashi, H., 1992. Slow inactivation of tetrodotoxin-insensitive Na⁺ channels in neurons of rat dorsal root ganglia. *J. Membr. Biol.* 129, 71–80.
- Oweiss, K.G., 2010. *Statistical Signal Processing For Neuroscience and Neurotechnology*. Academic Press.
- Pan, W.J., Thompson, G.J., Magnuson, M.E., Jaeger, D., Keilholz, S., 2013. Infralow LFP correlates to resting-state fMRI BOLD signals. *Neuroimage* 74, 288–297. doi:10.1016/j.neuroimage.2013.02.035.
- Pan, W.J., Thompson, G., Magnuson, M., Majeed, W., Jaeger, D., Keilholz, S., 2011. Broad-band local field potentials correlate with spontaneous fluctuations in functional magnetic resonance imaging signals in the rat somatosensory cortex under isoflurane anesthesia. *Brain Connect.* 1, 119–131. doi:10.1089/brain.2011.0014.
- Peer, M., Pruss, H., Ben-Dayyan, I., Paul, F., Arzy, S., Finke, C., 2017. Functional connectivity of large-scale brain networks in patients with anti-NMDA receptor encephalitis: an observational study. *Lancet Psychiatry* 4, 768–774. doi:10.1016/S2215-0366(17)30330-9.
- Power, J.D., Mitra, A., Laumann, T.O., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2014. Methods to detect, characterize, and remove motion artifact in resting state fMRI. *Neuroimage* 84, 320–341. doi:10.1016/j.neuroimage.2013.08.048.
- Power, J.D., Plitt, M., Gotts, S.J., Kundu, P., Voon, V., Bandettini, P.A., Martin, A., 2018. Ridding fMRI data of motion-related influences: removal of signals with distinct spatial and physical bases in multiecho data. *Proc. Natl. Acad. Sci. U. S. A.* 115, E2105–E2114. doi:10.1073/pnas.1720985115.
- 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. *Proc. Natl. Acad. Sci. U. S. A.* 98, 676–682. doi:10.1073/pnas.98.2.676.
- Salvador, R., Martínez, A., Pomarol-Clotet, E., Gomar, J., Vila, F., Sarró, S., Capdevila, A., Bullmore, E., 2008. A simple view of the brain through a frequency-specific functional connectivity measure. *Neuroimage* 39, 279–289. doi:10.1016/j.neuroimage.2007.08.018.
- Sasai, S., Homae, F., Watanabe, H., Taga, G., 2011. Frequency-specific functional connectivity in the brain during resting state revealed by NIRS. *Neuroimage* 56, 252–257. doi:10.1016/j.neuroimage.2010.12.075.
- Schmidt, S.L., Iyengar, A.K., Foulser, A.A., Boyle, M.R., Fröhlich, F., 2014. Endogenous cortical oscillations constrain neuromodulation by weak electric fields. *Brain Stimul.* 7, 878–889. doi:10.1016/j.brs.2014.07.033.
- Schölvinck, M.L., Maier, A., Ye, F.Q., Duyn, J.H., Leopold, D.A., 2010. Neural basis of global resting-state fMRI activity. *Proc. Natl. Acad. Sci. U.S.A.* 107, 10238–10243. doi:10.1073/pnas.0913110107.
- Schulz, K., Sydekum, E., Krueppel, R., Engelbrecht, C.J., Schlegel, F., Schroter, A., Rudin, M., Helmchen, F., 2012. Simultaneous BOLD fMRI and fiber-optic calcium recording in rat neocortex. *Nat. Methods* 9, 597–602. doi:10.1038/nmeth.2013.
- Shmuel, A., Leopold, D.A., 2008. Neuronal correlates of spontaneous fluctuations in fMRI signals in monkey visual cortex: implications for functional connectivity at rest. *Hum. Brain Mapp.* 29, 751–761. doi:10.1002/hbm.20580.
- Thompson, G.J., Merritt, M.D., Pan, W.J., Magnuson, M.E., Grooms, J.K., Jaeger, D., Keilholz, S.D., 2013. Neural correlates of time-varying functional connectivity in the rat. *Neuroimage* 83, 826–836. doi:10.1016/j.neuroimage.2013.07.036.
- Thompson, J.K., Peterson, M.R., Freeman, R.D., 2003. Single-neuron activity and tissue oxygenation in the cerebral cortex. *Science* 299, 1070–1072. doi:10.1126/science.1079220.
- Toib, A., Lyakhov, V., Marom, S., 1998. Interaction between duration of activity and time course of recovery from slow inactivation in mammalian brain Na⁺ channels. *J. Neurosci. Off. J. Soc. Neurosci.* 18, 1893–1903.
- Tong, Y., Hocke, L.M., Fan, X., Janes, A.C., Frederick, B., 2015. Can apparent resting state connectivity arise from systemic fluctuations? *Front. Hum. Neurosci.* 9, 285. doi:10.3389/fnhum.2015.00285.
- Vazquez, A.L., Murphy, M.C., Kim, S.G., 2014. Neuronal and physiological correlation to hemodynamic resting-state fluctuations in health and disease. *Brain Connect.* 4, 727–740. doi:10.1089/brain.2014.0276.
- Vincent, J.L., Patel, G.H., Fox, M.D., Snyder, A.Z., Baker, J.T., Van Essen, D.C., Zempel, J.M., Snyder, L.H., Corbetta, M., Raichle, M.E., 2007. Intrinsic functional architecture in the anaesthetized monkey brain. *Nature* 447, 83–86. doi:10.1038/nature05758.
- Wang, M., He, Y., Sejnowski, T.J., Yu, X., 2018. Brain-state dependent astrocytic Ca²⁺ signals are coupled to both positive and negative BOLD-fMRI signals. In: *Proceedings of the National Academy of Sciences*.
- Wilson, C., 2008. Up and down states. *Scholarpedia* J. 3, 1410.
- Khine, M., Thompson, J.K., Bierer, S., Altamirano, L.J., Freeman, R., and Lee, L.P., 2003. A combined neural-oxygen sensing multi-channel electrode array. *First International IEEE EMBS Conference on Neural Engineering*, 2003. Conference Proceedings, 204–207.