

Beringia 2005 expeditions funded by the Polar Research Secretariat at the Royal Swedish Academy of Sciences. The main work was funded by grant 150322/720 to Brochmann from the Research Council of Norway. Additional grants to Westergaard were obtained from K. and H. Jakobsens Fund, King Haakon VII Educational

Fund, Roald Amundsen's Centre for Arctic Research, KOMETEN, and Tromsø University Museum.

Supporting Online Material

www.sciencemag.org/cgi/content/full/316/5831/1606/DC1
Materials and Methods

Tables S1 and S2
References

22 December 2006; accepted 2 May 2007
10.1126/science.1139178

Modulation of Neuronal Interactions Through Neuronal Synchronization

Thilo Womelsdorf,^{1*†} Jan-Mathijs Schoffelen,^{1*†} Robert Oostenveld,¹ Wolf Singer,^{2,3} Robert Desimone,^{4,5} Andreas K. Engel,⁶ Pascal Fries^{1,7}

Brain processing depends on the interactions between neuronal groups. Those interactions are governed by the pattern of anatomical connections and by yet unknown mechanisms that modulate the effective strength of a given connection. We found that the mutual influence among neuronal groups depends on the phase relation between rhythmic activities within the groups. Phase relations supporting interactions between the groups preceded those interactions by a few milliseconds, consistent with a mechanistic role. These effects were specific in time, frequency, and space, and we therefore propose that the pattern of synchronization flexibly determines the pattern of neuronal interactions.

Groups of activated neurons synchronize in the gamma-frequency band (30 to 100 Hz), and previous studies have related gamma-band synchronization to several cognitive functions (1–6). Yet, if gamma-band synchronization subserves those functions, it must have mechanistic consequences for neuronal processing (7). It has been shown that the precise timing of pre- and postsynaptic activation determines long-term changes in synaptic strength (8–10) and that gamma-band synchronization of synaptic inputs directly enhances their effective synaptic strength (11–13).

Synchronization between two groups of neurons is also likely to facilitate interactions between them (Fig. 1A) (6, 14). Gamma-band synchronization entails rhythmic inhibition of the local network (15–17), and the periods between inhibition provide temporal windows for neuronal interaction. Two groups of neurons will therefore probably have a greater influence on each other when their temporal interaction windows open at the same times, i.e., when the

rhythmic synchronization within the groups is also synchronized between the groups. By the same token, the interaction is probably curtailed if the temporal interaction windows open either in an uncorrelated way or consistently out of phase with each other.

We analyzed four data sets: (i) one from awake cat area 17, (ii) one combining awake cat area 18 with area 21a recordings, (iii) one from awake monkey area V1, and (iv) one from monkey area V4. [Data from two of the three area 17 data sets have been used in (18, 19); the V4 data set has been used in (3, 20).] In all cases, we recorded multiunit activity (MUA) and local field potentials (LFPs) simultaneously from four to eight electrodes while the neurons were visually stimulated with moving gratings. From each data set, we used trials with identical visual stimulation and behavioral tasks and based our analysis on the natural fluctuation of neuronal gamma-band synchronization. For each pair of neuronal groups, we quantified synchronization by means of the MUA-MUA phase-coherence spectrum (Fig. 1B) and the MUA-LFP phase-coherence spectrum (Fig. 1C) (21).

Phase-coherence spectra showed a peak in the gamma-frequency band, indicating that phase relations between signals were not random. However, phase coherence was far from perfect (a value of 1.0), but it assumed average peak values of 0.14 and 0.27 for MUA-MUA and MUA-LFP combinations, respectively. The phase relations at 60 Hz in one example MUA-MUA pair are shown for 708 trials of 250-ms length (phase-coherence value of 0.06) (Fig. 1D).

The spread of phase relations around their mean might just be irrelevant noise. Here, however, we used this spread to actually test for its potential physiological consequences. We hypothesized that the mutual influence between two

neuronal groups was a function of their phase relation (Fig. 1A). Phase relations are meaningfully defined per frequency, and we hypothesized that the phase relation at a given frequency should modulate the interaction among the local rhythmic activities specifically at that frequency.

We investigated this hypothesis for the example pair of recordings sites. We sorted the trials into six bins according to the 60-Hz phase relation between the two MUAs (Fig. 1D). For each phase-relation bin separately, we then quantified the two MUAs' mutual influence as the Spearman rank correlation coefficient between the two MUAs' 60-Hz power, across the trials in the bin (Fig. 1E). Fluctuations of 60-Hz power were most strongly correlated when the 60-Hz phase relation was close to its mean across the trials. Specifically, when the gamma-band rhythm in group A led the one in group B by 2.1 ms (mean phase relation at 45.8°), the correlation between each group's gamma-band power was four times as strong as when the rhythms were separated by 10.5 ms (phase relation at 225.8°). The example pair illustrates this for a case with a nonzero mean phase to demonstrate that the effect cannot be ascribed to external artifacts or volume conduction, but the mean phase relations across our sample distributed closely around zero (Fig. 1B).

We performed the same analysis after replacing one of the MUAs by the LFP recorded through the same electrode. The mean MUA-LFP phase relations clustered around 141° (Fig. 1C), and power correlations were again substantially enhanced around the mean phase relation (Fig. 1, F and G). Across our sample, good phase relations mostly distributed close to the respective mean phase relations for both MUA-MUA and MUA-LFP pairs (fig. S1). We correspondingly dubbed the mean phase relation as “good” and the opposite phase relation as “bad,” and we aligned the trial binning to the good phase relation.

The observed effect was consistent across the four data sets (Fig. 2 and fig. S2) (140, 86, 111, and 111 MUA-MUA pairs from area 17, areas 18×21a, area V1, and area V4, respectively, and 280, 172, 228, and 237 MUA-LFP pairs from the same areas). MUA-LFP pairs showed qualitatively the same effect as MUA-MUA pairs but with higher signal-to-noise ratios (Fig. 2B). We therefore focused our further analyses on MUA-LFP pairs (recorded from separate electrodes). The effect was also present for pairs of LFP and single-unit recordings (fig. S3). The effect generalized to long-range interactions, because the analysis of the data set combining cat area

¹F. C. Donders Centre for Cognitive Neuroimaging, Radboud University Nijmegen, 6525 EN Nijmegen, Netherlands.

²Department of Neurophysiology, Max Planck Institute for Brain Research, 60528 Frankfurt, Germany. ³Frankfurt Institute for Advanced Studies, Johann Wolfgang Goethe University, 60438 Frankfurt, Germany. ⁴Laboratory of Neuropsychology, National Institute of Mental Health, National Institutes of Health, Bethesda, MD 20892, USA.

⁵McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.

⁶Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf, 20246 Hamburg, Germany. ⁷Department of Biophysics, Radboud University Nijmegen, 6525 EZ Nijmegen, Netherlands.

*These authors contributed equally to this work.

†To whom correspondence should be addressed. E-mail: thilo.womelsdorf@fcdonders.ru.nl (T.W.); jan.schoffelen@fcdonders.ru.nl (J.-M.S.)

18 recordings with area 21a recordings was restricted to interareal pairs of recording sites (Fig. 2D).

We did not attempt to relate the two signals' power correlation to their coherence, because

the coherence measure confounds phase synchronization and power correlation. In contrast, we related the correlation between the two signals' power across trials directly to their relative phase. The presence of coherence (Fig. 1H)

does not necessarily result in a phase-relation-dependent power correlation (Fig. 1I) and thus, the demonstration of phase-relation-dependent power correlation goes beyond the demonstration of coherence.

Fig. 1. Precise timing between rhythmic neuronal activities determines the strength of their mutual influence. (A) Sketch of three groups of neurons, each rhythmically active (LFP oscillations with spikes in troughs). Time windows for effective communication are either aligned (red and blue group) or not aligned (red and gray group). (B and C) Average phase-coherence spectrum across all (B) MUA-MUA and (C) MUA-LFP pairs (area 17 data) and corresponding distributions of mean phase relations at 60 Hz. (D) Trialwise phase relations from an example MUA-MUA pair. Phase relations were sorted into bins (light and dark gray ring segments) aligned to the mean phase relation (red line). (E) Spearman rank correlation coefficients between the two MUAs' 60-Hz power as a function of their phase relation. (The solid line indicates a cosine fit.) (F and G) Same as (D) and (E), but with one MUA substituted by the respective LFP. (H and I) Example MUA-LFP pair from the area 17 data set demonstrating that coherence does not necessarily result in phase-relation-dependent power correlations. (H) Coherence with a clear peak around 60 Hz. (I) Power correlations as a function of phase relations, showing no consistent relation.






Fig. 2. Phase-relation-dependent modulation of power correlations is frequency specific. (A) Average power correlation as a function of phase relation (x axis) and frequency (y axis) for MUA-MUA pairs recorded in cat area 17. (B) Same as (A), but for MUA-LFP pairs. (C) Modulation depth of the cosine function

fitted to the phase-relation-dependent power correlations. Gray bars indicate significant frequencies ($P < 0.05$, multiple comparisons corrected). (Right) Average phase-relation-dependent power correlation at 60 Hz. (D and E) Same as (C), but for (D) cat area 18x21a and (E) monkey area V1.

One concern is that the effect might be due to rhythmic common input that fluctuates in strength across trials. Common input would bias phase relations and impose correlated power on both groups. Clues about the actual causal chain of events might be gained from the relative timing between phase relation, on one hand, and power correlation, on the other hand. If power correlation covaried with phase relation because of common input, then there should be no delay between the two. However, if good phase relations were actually the mechanistic cause of strong power correlations, then good phase relations should precede strong power correlations by a few milliseconds, incurred by axonal, synaptic, and intracellular delays. We therefore compiled time-resolved estimates of the “goodness” of phase relations and of the strength of power correlations, and then we determined the cross-correlation as a function of time lag between the two time series (21). Figure 3 shows this analysis pooled across all four data sets and demonstrates that good phase relations preceded strong power correlations by 5 ms. We

observed this temporal precedence in each of the four data sets. Additional observations arguing against a common input explanation are given in the supporting online material text.

The modulation of two neuronal groups’ interaction by their phase relation would be particularly interesting if it was spatially specific (Fig. 1A). We therefore investigated recording triplets A, B, and C from three separate electrodes (in which A was an LFP and B and C were MUAs). Figure 4A shows, for one example triplet, a scatter plot in which each dot corresponds to one trial and the *x* and *y* values give the “goodness” of the A-B and A-C phase relations, respectively (22). We sorted the trials according to the quadrants of the plot. In quadrants Q1 and Q2, the phase relation between A and B was good, whereas in Q3 and Q4 it was bad. We contrasted the A-B power correlation for Q1 and Q2 with that for Q3 and Q4 (red line in Fig. 4, B to D). Orthogonally to this, the phase relation between A and C was good in Q1 and Q3, whereas it was bad in Q2 and Q4, and we contrasted the A-B power correlation for Q1

and Q3 with that for Q2 and Q4 (blue line in Fig. 4, B to D). The A-B power correlation depended significantly more on the A-B phase relation than on the A-C phase relation. Thus, the effect had a spatial resolution that was at least as high as the spatial resolution of our recordings (down to 0.65 mm in the monkey data sets).

We provided evidence suggesting that neuronal interactions mechanistically depend on the phase relation between rhythmic activities. The most likely reason for this dependence is that rhythmic activities modulate the gain of incoming synaptic input rhythmically. Effective connectivity can thus be maximized or minimized through synchronization at a good or bad phase relation. The impact of pyramidal cells could be enhanced, for example, if their firing phase relative to interneurons were advanced (15, 23, 24) or if interneuronal firing were delayed through inhibition or reduced excitation (25). Such mechanisms might be invoked directly by cognitive top-down control.

Effective connectivity would diminish when synchronization is less precise, because then synaptic input is more likely to arrive at random phases. This mechanism has the advantage that within a sufficiently wide frequency band, multiple groups can be desynchronized, with respect to a given target group, without being necessarily synchronized to each other. Periods of putative interactions between distant neuronal groups are marked by an increased precision of synchronization (1, 4, 6, 26–30).

We propose that the pattern of synchronization (its precision, phase, or both) weights the anatomical-connection infrastructure with a gain pattern, resulting in an effective interaction pattern (14). Such a mechanism would have several





Fig. 4. Spatial selectivity of phase-relation-dependent power correlation. **(A)** Scatterplot shows the distribution of trialwise phase relations between groups A and B (*y* axis) and between groups A and C (*x* axis) for an example triplet at 60 Hz. Equations define how A-B power correlations from each quadrant were combined for the results shown in (B to D). In the equations, $c(AB_q)$ denotes the A-B power correlation across trials in quadrant *q* (where *q* is 1, 2, 3, or 4). **(B)** A-B power correlation as a function of the A-B phase relation [irrespective of the A-C phase relation (red line)] and as a function of the A-C phase relation [irrespective of the A-B phase relation (blue line)]. Gray bars indicate frequencies with significant differences (*P* < 0.05, multiple comparisons corrected). The *y* axis denotes the differences in power correlations according to the equations shown in (A). **(C and D)** Same as (B), but for (C) monkey area V1 and (D) monkey area V4.



interesting features. First, the effective interaction pattern could be modified very dynamically. Second, the mechanism would act connection-wise. Third, a transient interaction would lead to spike-time-dependent plasticity (8–10) and thus, a long-term trace. And fourth, synchronization might emerge in a self-organized manner between “matching” neuronal groups. In the visual cortex, synchronization is stronger among neurons activated by the same visual stimulus (*I*). This principle might generalize to the handshaking between cognitive top-down control and matching sensory bottom-up information, in which case consecutive synchronization could contribute to the selective routing of sensory information to behavioral control (13, 14, 25). Our results suggest that synchronization has consequences for neuronal interactions, providing a putative mechanism through which synchronization contributes to cognitive functions.

References and Notes

- C. M. Gray, P. König, A. K. Engel, W. Singer, *Nature* **338**, 334 (1989).
- A. Bragin *et al.*, *J. Neurosci.* **15**, 47 (1995).
- P. Fries, J. H. Reynolds, A. E. Rorie, R. Desimone, *Science* **291**, 1560 (2001).
- J. Fell *et al.*, *Nat. Neurosci.* **4**, 1259 (2001).
- B. Pesaran, J. S. Pezaris, M. Sahani, P. P. Mitra, R. A. Andersen, *Nat. Neurosci.* **5**, 805 (2002).
- J.-M. Schoffelen, R. Oostenveld, P. Fries, *Science* **308**, 111 (2005).
- G. Buzsáki, *Rhythms of the Brain* (Oxford Univ. Press, Oxford, 2006).
- H. Markram, J. Lübke, M. Frotscher, B. Sakmann, *Science* **275**, 213 (1997).
- G. Q. Bi, M. M. Poo, *J. Neurosci.* **18**, 10464 (1998).
- V. Wespapat, F. Tegnigkeit, W. Singer, *J. Neurosci.* **24**, 9067 (2004).
- P. König, A. K. Engel, W. Singer, *Trends Neurosci.* **19**, 130 (1996).
- R. Azouz, C. M. Gray, *Neuron* **37**, 513 (2003).
- E. Salinas, T. J. Sejnowski, *Nat. Rev. Neurosci.* **2**, 539 (2001).
- P. Fries, *Trends Cognit. Sci.* **9**, 474 (2005).
- A. Hasenstaub *et al.*, *Neuron* **47**, 423 (2005).
- I. Vida, M. Bartos, P. Jonas, *Neuron* **49**, 107 (2006).
- M. Bartos, I. Vida, P. Jonas, *Nat. Rev. Neurosci.* **8**, 45 (2007).
- P. Fries, P. R. Roelfsema, A. K. Engel, P. König, W. Singer, *Proc. Natl. Acad. Sci. U.S.A.* **94**, 12699 (1997).
- P. Fries, J. H. Schröder, P. R. Roelfsema, W. Singer, A. K. Engel, *J. Neurosci.* **22**, 3739 (2002).
- T. Womelsdorf, P. Fries, P. P. Mitra, R. Desimone, *Nature* **439**, 733 (2006).
- Materials and methods are available as supporting material on Science Online.
- Positive and negative departures from the good phase relation toward the bad phase relation were pooled because our previous analyses had demonstrated symmetric effects in both directions.
- P. König, A. K. Engel, P. R. Roelfsema, W. Singer, *Neural Comput.* **7**, 469 (1995).
- J. Csicsvari, B. Jamieson, K. D. Wise, G. Buzsáki, *Neuron* **37**, 311 (2003).
- J. Mishra, J. M. Fellous, T. J. Sejnowski, *Neural Networks* **19**, 1329 (2006).
- A. K. Engel, A. K. Kreiter, P. König, W. Singer, *Proc. Natl. Acad. Sci. U.S.A.* **88**, 6048 (1991).
- G. Buzsáki, Z. Horváth, R. Urioste, J. Hetke, K. Wise, *Science* **256**, 1025 (1992).
- P. R. Roelfsema, A. K. Engel, P. König, W. Singer, *Nature* **385**, 157 (1997).
- C. Tallon-Baudry, S. Mandon, W. A. Freiwald, A. K. Kreiter, *Cereb. Cortex* **14**, 713 (2004).
- J. Gross *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **101**, 13050 (2004).
- We thank P. König, P. R. Roelfsema, and J. H. Schröder for help during the cat experiments and J. H. Reynolds, A. E. Rorie, and A. F. Rossi for help during the monkey experiments. This work was supported by the Human Frontier Science Program and the Netherlands Organization for Scientific Research (P.F.); the Volkswagen Foundation (A.K.E. and P.F.); the Max Planck Society (P.F., A.K.E., and W.S.); European Union grants EU IST-027268 and EU NEST-043457 (A.K.E.); and NIH grant R01EY017292 and the National Institute of Mental Health—Intramural Research Program (R.D.).

Supporting Online Material

www.sciencemag.org/cgi/content/full/316/5831/1609/DC1
Materials and Methods
SOM Text
Figs. S1 to S6
References

5 January 2007; accepted 1 May 2007
10.1126/science.1139597

Neural Mechanisms of Visual Attention: How Top-Down Feedback Highlights Relevant Locations

Yuri B. Saalman,¹ Ivan N. Pigarev,² Trichur R. Vidyasagar^{1*}

Attention helps us process potentially important objects by selectively increasing the activity of sensory neurons that represent the relevant locations and features of our environment. This selection process requires top-down feedback about what is important in our environment. We investigated how parietal cortical output influences neural activity in early sensory areas. Neural recordings were made simultaneously from the posterior parietal cortex and an earlier area in the visual pathway, the medial temporal area, of macaques performing a visual matching task. When the monkey selectively attended to a location, the timing of activities in the two regions became synchronized, with the parietal cortex leading the medial temporal area. Parietal neurons may thus selectively increase activity in earlier sensory areas to enable focused spatial attention.

Attention allows us to engage with our environment by selecting information relevant for behavior (1–3). This enables preferential processing of particular locations in the visual field or specific features of objects. Attention maintained on a location is usually referred to as spatial attention and that on a

feature as feature-based attention (4, 5). Both types of attention manifest in visual cortical areas as increased activity of neurons representing the attended location or feature and reduced activity of other neurons (6–14). This may require top-down feedback about what is relevant in the environment; however, such feedback has not been empirically demonstrated. There is evidence (6, 11, 15–17) that the posterior parietal cortex (PPC) is critical for spatial attention. The PPC is a higher-order structure along the dorsal stream of visual areas (Fig. 1A), which are particularly concerned with spatial aspects of a scene. It has been suggested that the spatial information about

a scene extracted by the PPC forms the basis for feedback signals to earlier levels of the visual pathway, highlighting spatial locations of potential interest (18, 19) and gating responses depending upon the state of attention.

We simultaneously measured the activity in a part of the macaque PPC called the lateral intraparietal area (LIP), and the immediately earlier stage of the dorsal pathway, the medial temporal area (MT; Fig. 1A). We tested whether LIP feedback increases MT responses to attended visual stimuli. The monkeys performed a delayed match-to-sample (DMS) task, which manipulated both where they were attending and what stimulus feature they were attending to (Fig. 1B) (20). We recorded 29 pairs of neurons from MT and LIP, each pair having overlapping receptive fields (RFs) and the same preferred orientation. Local field potentials (LFPs) from the two sites were also recorded.

We first tested whether our paradigm resulted in increased responses of MT neurons to attended stimuli. Figure 2A shows the effect of spatial attention on a single MT neuron and Fig. 2B, the average population response. For both data sets, the MT response to the second stimulus was significantly increased in the “spatial and feature-based attention” and “spatial attention” conditions, compared to the “neutral” control. When “attention was elsewhere,” the MT response to the second stimulus was significantly reduced (Holm’s controlled Wilcoxon test, $P < 0.05$). These attentional effects on MT neurons are consistent with those reported in other types of cognitive

¹Department of Optometry and Vision Sciences, The University of Melbourne, Parkville 3010, Australia. ²Institute for Information Transmission Problems, Russian Academy of Sciences, Bol’shoy Karetniy 19, 127994 Moscow, Russia.

*To whom correspondence should be addressed. E-mail: trv@unimelb.edu.au