ARTICLE

Spontaneous variability in gamma dynamics described by a damped harmonic oscillator driven by noise

Georgios Spyropoulos [1✉], Matteo Saponati[1,2], Jarrod Robert Dowdall[1,2], Marieke Louise Schölvinck[1], Conrado Arturo Bosman[3,4], Bruss Lima [5,6], Alina Peter [1,2], Irene Onorato[1,2], Johanna Klon-Lipok[1,5], Rasmus Roese[1], Sergio Neuenschwander[5,7], Pascal Fries [1,3,9✉] & Martin Vinck [1,8,9✉]

Circuits of excitatory and inhibitory neurons generate gamma-rhythmic activity (30–80 Hz). Gamma-cycles show spontaneous variability in amplitude and duration. To investigate the mechanisms underlying this variability, we recorded local-field-potentials (LFPs) and spikes from awake macaque V1. We developed a noise-robust method to detect gamma-cycle amplitudes and durations, which showed a weak but positive correlation. This correlation, and the joint amplitude-duration distribution, is well reproduced by a noise-driven damped harmonic oscillator. This model accurately fits LFP power-spectra, is equivalent to a linear, noise-driven E-I circuit, and recapitulates two additional features of gamma: (1) Amplitude-duration correlations decrease with oscillation strength; (2) amplitudes and durations exhibit strong and weak autocorrelations, respectively, depending on oscillation strength. Finally, longer gamma-cycles are associated with stronger spike-synchrony, but lower spike-rates in both (putative) excitatory and inhibitory neurons. In sum, V1 gamma-dynamics are well described by the simplest possible model of gamma: A damped harmonic oscillator driven by noise.

[1] Ernst Strüngmann Institute (ESI) for Neuroscience in Cooperation with Max Planck Society, 60528 Frankfurt, Germany. [2] International Max Planck Research School for Neural Circuits, Frankfurt Am Main, Germany. [3] Donders Institute for Brain, Cognition and Behaviour, Radboud University, 6525 EN Nijmegen, the Netherlands. [4] Swammerdam Institute for Life Sciences, Center for Neuroscience, Faculty of Science, University of Amsterdam, 1098 XH Amsterdam, the Netherlands. [5] Max Planck Institute for Brain Research, 60438 Frankfurt, Germany. [6] Carlos Chagas Filho Institute of Biophysics, Federal University of Rio de Janeiro, 21941-902 Rio de Janeiro, Brazil. [7] Brain Institute, Federal University of Rio Grande do Norte, 59056-450 Natal, Brazil. [8] Donders Centre for Neuroscience, Department of Neuroinformatics, Radboud University, 6525 EN Nijmegen, the Netherlands. [9] These authors contributed equally: Pascal Fries, Martin Vinck. ✉email: georgios.spyropoulos@esi-frankfurt.de; pascal.fries@esi-frankfurt.de; martin.vinck@esi-frankfurt.de

The brain consists of different kinds of cell types, which have unique properties and are commonly divided into inhibitory (I) and excitatory (E) neurons. E-I Interactions can generate collective rhythmic activity in different frequency bands. One of the "faster" rhythms in neocortical circuits is the gamma rhythm (30–80 Hz), whose function has been heavily debated in the literature[1–15]. This rhythm can be observed at many scales, from the macro/meso-scale (MEG, EEG, ECoG, LFP), to the microscale (synaptic currents and spiking activity)[6,16,17]. It is however unknown how the properties of collective neuronal gamma synchronization can arise from interactions between its microscopic constituents[16,18].

Analysis of macro/meso-scopic gamma dynamics has revealed substantial variability in the amplitude and frequency of gamma oscillations as a function of time, but also cortical space[4,11,19–23]. In particular, gamma oscillations are not well approximated by sinusoids[4], despite the fact that they are often depicted as such. Rather, they show major fluctuations in amplitude over time, sometimes described as "bursts"; as well as their frequency, giving rise to the broad-band spectral nature of gamma. These fluctuations likely reflect the properties of the underlying E-I circuit and the way it responds to changes in input drive, and they are relevant for the possible functional roles of gamma[5,8,10,11,13,23]. Previous work in rodent hippocampus[24] has suggested that cycle-by-cycle fluctuations in amplitude and duration (i.e. the inverse of frequency) are explained by two model components: (1) cycle-by-cycle fluctuations in synaptic excitation; and (2) balanced, bidirectional interactions between E and I neurons, consistent with the PING (Pyramidal Interneuronal Network Gamma) model of the gamma rhythm[3,6,25–29]. The proposed model for hippocampus holds that the occurrence of a strong bout of synaptic excitation will be balanced by high-amplitude, long-lasting inhibition. As predicted from this model, gamma-cycle amplitude and duration were reported to be strongly correlated ($r = 0.61$) in rodent hippocampus[24].

The starting point of the present study was to see whether this regularity generalizes to other cortical circuits, in particular to awake primate visual cortex, another system where gamma oscillations have been extensively studied. It remains unclear how the mechanisms of gamma in visual cortex compare to hippocampus. It appears that E-I mechanisms of gamma in higher visual areas (V4) might be comparable to hippocampus[29], although there is evidence that they are substantially different in primary visual cortex (V1)[30]. Furthermore, the dependence of V1/V2 gamma on stimulus contrast suggests that increases in synaptic excitation lead to *increases* rather than *decreases* in the frequency of V1/V2 gamma[20,31]. It is unknown, however, what the relationship is between *spontaneous* fluctuations in gamma-cycle amplitude and duration in V1. Our paper consists of two parts: In the first part (Figs. 1–5), we show a positive correlation between spontaneous fluctuations in gamma-cycle amplitude and duration and address several confounds inherent to cycle-by-cycle analyses. In the second part (Figs. 6–9), we examine the mechanisms underlying these correlations by analyzing spiking activity and drawing a comparison with a damped harmonic oscillator model driven by noise.

Results

Recordings and task. We recorded LFPs and spiking activity from area V1 of several awake macaque monkeys (see Methods). Monkeys performed a fixation task while drifting gratings or uniform colored surfaces were presented. Figure 1a shows an example trial of broad-band LFP recorded during the presentation of a full-screen drifting grating. The trial-average spectra of absolute power (Fig. 1b) and of the power-change relative to pre-

stimulus baseline (Fig. 1c) reveal strong visually-induced gamma oscillations. Time-frequency analysis (Fig. 1d) shows that this induced gamma rhythm is sustained for the duration of the visual-stimulation period. Figure 1f–i shows similar results for colored surface stimuli[32,33].

The correlation between gamma-cycle amplitude and duration. A previous study has examined correlations between the amplitude and duration of individual gamma cycles in rat hippocampus CA3[31], and reported a strongly positive ($r = 0.61$) correlation, both in vivo and in vitro. We wondered whether a similarly strong correlation exists in monkey V1. We therefore used the same analysis method as previously used for rat hippocampus. This method is based on (1) band-pass filtering LFP signals, (2) detecting periods of high-amplitude gamma activity, and (3) detecting empirical peaks and troughs in the filtered signal (Fig. 2a, b; see Methods section). Using this method, we found a relatively strong positive ($r = 0.361$) correlation between the amplitude and duration of individual gamma cycles during the visual stimulation period (Fig. 2c). By contrast, correlations between the amplitude of a given cycle and the duration of either the preceding or succeeding cycle were not significant (Fig. 2c).

We expected that this result would be specific to the visual stimulation period, during which gamma oscillations were prominent, but that it would not hold for the pre-stimulus period, during which there was no visible gamma peak in the LFP power spectrum (Fig. 1b, g). Yet, for the pre-stimulus period, the algorithm detailed above detected a substantial amount of gamma epochs. Surprisingly, we observed even stronger correlations between gamma-cycle amplitudes and durations for the pre-stimulus ($r = 0.605$) compared to the stimulus period (Fig. 2d).

This prompted us to investigate whether the same algorithm would also detect a positive correlation between gamma-cycle amplitudes and durations for synthetic $1/f^n$ noise signals (Fig. 2e, f). This was indeed the case (Fig. 2g). Thus, noisy fluctuations in a signal without rhythmic components can give rise to a strong positive correlation between the amplitudes and durations of detected "gamma cycles". The presence of a positive correlation between amplitudes and durations can be made intuitive by considering a random walk process: In such a process, the magnitudes of successive increments or decrements are independent of each other, with zero mean. In this case, a successive series of positive increments typically results in a "cycle" with a high amplitude and a long duration. By contrast, a rapid reversal typically results in a low-amplitude "cycle" with a short duration. Hence, the positive correlation between gamma-cycle amplitude and duration in the stimulus period may have been due to noisy background fluctuations instead of oscillatory activity.

This prompted us to develop a method that (1) avoided band-pass filtering in a narrow frequency-range; and (2) ensured that gamma peaks and troughs were not detected due to noisy fluctuations, but reflected a rhythmic process (Fig. 3a–d; see Methods section). Note that in our datasets, only gamma-cycles fulfilled the necessary criteria on oscillatory behavior as defined by the Hilbert phase spectrum. To obtain estimates of gamma-cycle amplitudes and durations with a high temporal resolution, we measured them in periods of "half-cycles" (i.e. peak-to-trough or trough-or-peak). (For the rest of the text, we will be referring to the amplitudes and durations of individual gamma half-cycles as "gamma-cycle amplitudes" and "gamma-cycle durations", and will mention explicitly when we measure them in full rather than half cycles). In contrast to the previous method (Fig. 2), our method detected very few gamma cycles in the pre-stimulus period (Fig. 3a–c). Because of this, a correlation between gamma-cycle amplitude and duration could not be reliably computed for this
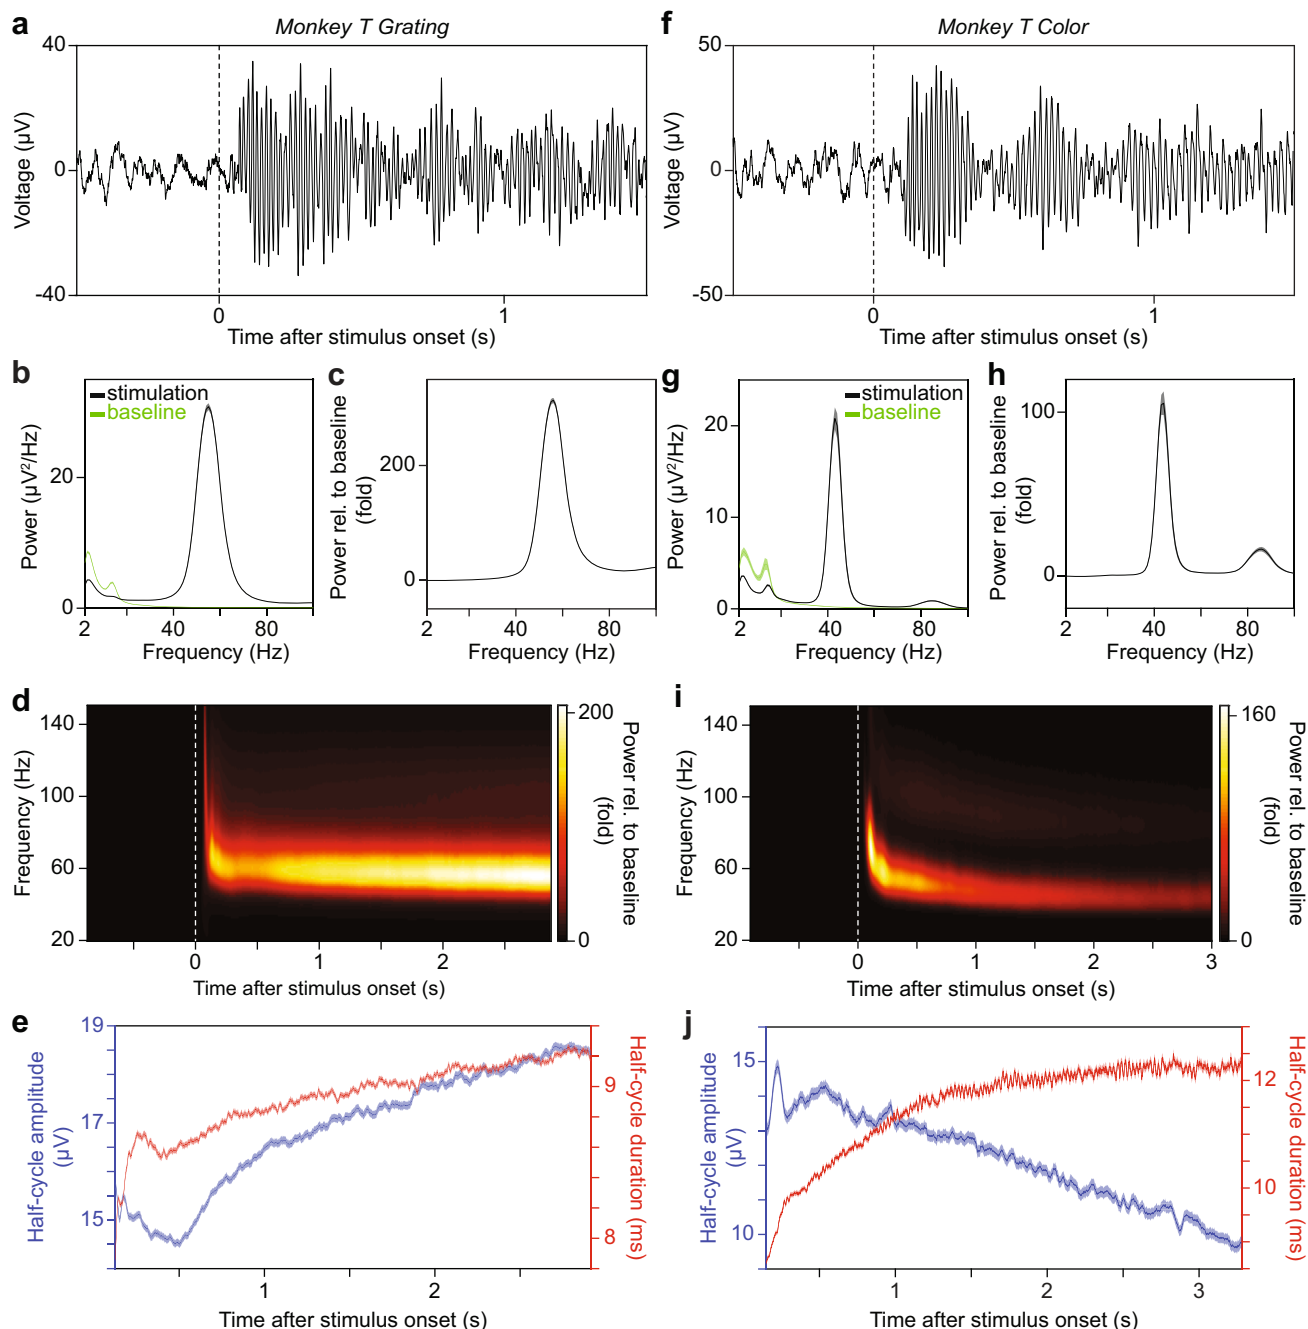
Fig. 1 Gamma dynamics in awake macaque V1 during visual stimulation. a Raw LFP trace from one representative recording site from area V1 in monkey T before and during the presentation of a full-screen drifting grating. **b**, **c** Raw power (**b**) and power-change relative to baseline (**c**), averaged across all selected recording sites from V1 in monkey T. The green and black traces in **b** correspond to the pre-stimulus baseline period and stimulation period respectively. The error regions show 2 standard errors of the mean (S.E.M.) based on a bootstrap procedure across trials (1000 bootstraps). **d** Power change relative to baseline, as function of frequency and time relative to stimulus onset, averaged over all selected V1 recording sites in monkey T before and during the presentation of a full-screen drifting grating. Note the changes in gamma amplitude and frequency with time after stimulus onset. **e** Time course of gamma-half-cycle amplitude (blue) and duration (red), averaged over all selected V1 recording sites in monkey T during the presentation of a full-screen drifting grating. The error regions show ±2 SEM based on a bootstrap procedure. Only the stimulation period is shown, because only very few gamma cycles of very low amplitude were detected before stimulus onset. **f–j** Same as **a–e**, but for the presentation of a full-screen uniform color surface. **a**, **d**, **f**, **i** Dashed lines indicate stimulus onset.

period. To further examine the noise-robustness of our method, we simulated an AR(2) (2nd order auto-regressive) process that had a positive correlation between gamma-cycle amplitudes and durations in the absence of noise. We then added $1/f^2$ background noise of different intensities (see Methods section). We found that our method did not yield spurious correlations due to the

inclusion of noise; instead it failed to detect any gamma cycles for higher noise-levels (red line in Fig. 3e). By contrast, the previous method (Fig. 2) produced higher correlations when the noise-level increased (black line in Fig. 3e).

Using this new method, we then detected gamma-cycle amplitudes and durations for all trials and available time-points,
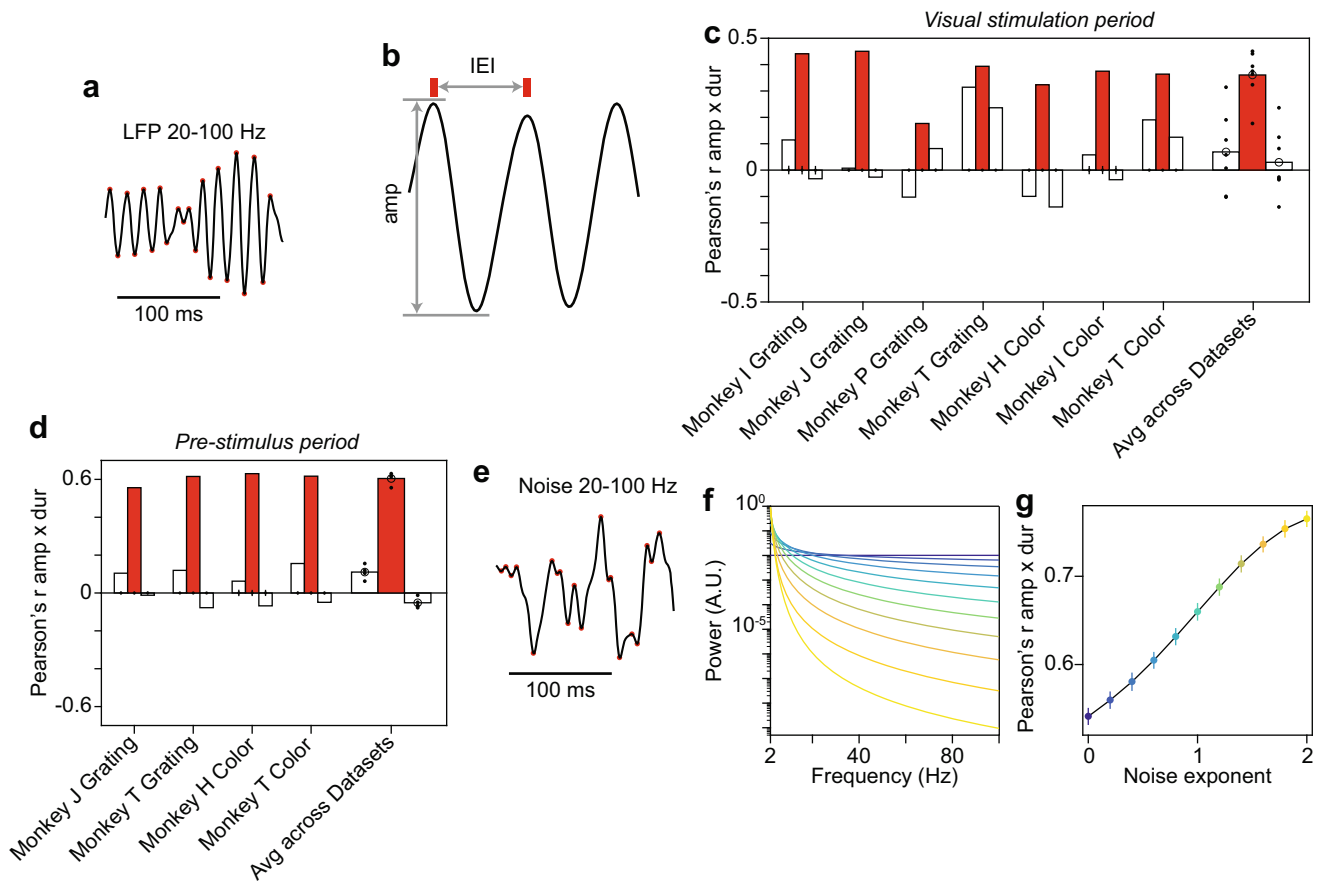
Fig. 2 Estimation of correlation between gamma-cycle amplitude and duration can be influenced by noise. a LFP trace filtered in the gamma range (20–100 Hz). Red dots indicate local maxima and minima. **b** Segment of the trace in **a** demonstrating the definition of gamma-cycle amplitude and inter-event interval, i.e. gamma-cycle duration. **c** For each dataset listed on the x-axis, the three bars show the correlation between gamma-cycle amplitudes and the durations of the same gamma cycle (center, red), the previous gamma cycle (left, white), and the next gamma cycle (right, white). On the right, this is shown for the average across all datasets. This was calculated for the visual stimulation period. Amplitude and duration values were extracted as in Roberts et al.[31], including the filtering illustrated in **a**, **b**; note that the employed subtraction of a boxcar-smoothed signal amounts to a high-pass-filtering at 20 Hz. For each dataset, a null distribution was produced by randomizing the order of duration values across trials, and the resulting means and 99.9% confidence intervals are shown as dots and vertical lines (numbers of trials = 278, 5740, 672, 1075, 672, 320, 142 for respective datasets). For the average across datasets, shown on the right, we performed a two-sided t-test and show the resulting confidence intervals as vertical lines on the observed mean (red bar: $p < 5 \times 10^{-5}$, white bars for preceding cycle: $p = 0.28$, white bars for succeeding cycle: $p = 0.56$). In addition, we performed a two-sided non-parametric permutation test (red bar: $p < 0.05$; white bars: $p > 0.05$). $n = 6$ biologically independent experiments. **d** Same as **c** but for the pre-stimulus baseline (numbers of trials = 5740, 1075, 672, 142 for respective datasets; averages across datasets: red bar: $p = 4.51 \times 10^{-5}$, two-sided t-test across datasets; white bars $p = 0.011$ and $p = 0.038$, respectively for preceding and succeeding cycles, two-sided t-test across datasets). $n = 6$ biologically independent experiments. **e** Example synthetic colored noise trace filtered in the gamma range (20–100 Hz). Red dots indicate local maxima and minima. **f** Power spectra of synthetic colored noise signals with a spectral shape of $1/f^n$, with n assuming values from 0 (dark blue) to 2 (bright yellow). **g** Correlation of the amplitude and duration of individual deflections in synthetic colored noise signals. Dots and vertical lines indicate means ±2 SEM produced by a bootstrap procedure (1000 bootstraps). The color conventions are the same as in **f**. Source data are provided in the Source Data file.

separately for each channel and stimulus condition. Because we were interested in spontaneous variability, we further ensured that correlations between gamma- cycle amplitudes and durations could not arise due to the time courses of amplitude and duration after stimulus onset (Fig. 1e, j). We achieved this by computing correlations across trials, separately for each available post-stimulus time-point and then averaging the correlations over time-points (see Methods section). With this approach, we found that the amplitudes and durations of individual gamma half-cycles were positively correlated in all tested datasets (Fig. 4a). The magnitude of these correlations was, on average, substantially lower (rho = 0.199; Pearson's r = 0.197) than the one observed with the previously employed method (compare Figs. 2c and 4a, t-test between Pearson's r across datasets, $p = 0.0084$). In addition, we computed the correlation between the amplitude

of a given half-cycle and the duration of the previous or the subsequent half-cycles, and this did not result in a consistent pattern of correlations across datasets (white bars in Fig. 4a). Similar results were obtained for full rather than half cycles, with significant correlations only for the same cycle comparison, but not for the preceding and succeeding cycle (Supplementary Fig. 1a).

The influence of slow dynamics and microsaccades. We wondered whether the observed correlation between gamma-cycle amplitudes and durations may have resulted from correlated changes in amplitudes and durations at relatively slow time-scales, e.g. due to drifts or slow oscillations in the monkey's state, or stimulus repetition effects[33–35]. In order to control for the potential influence of such changes, we computed the correlation
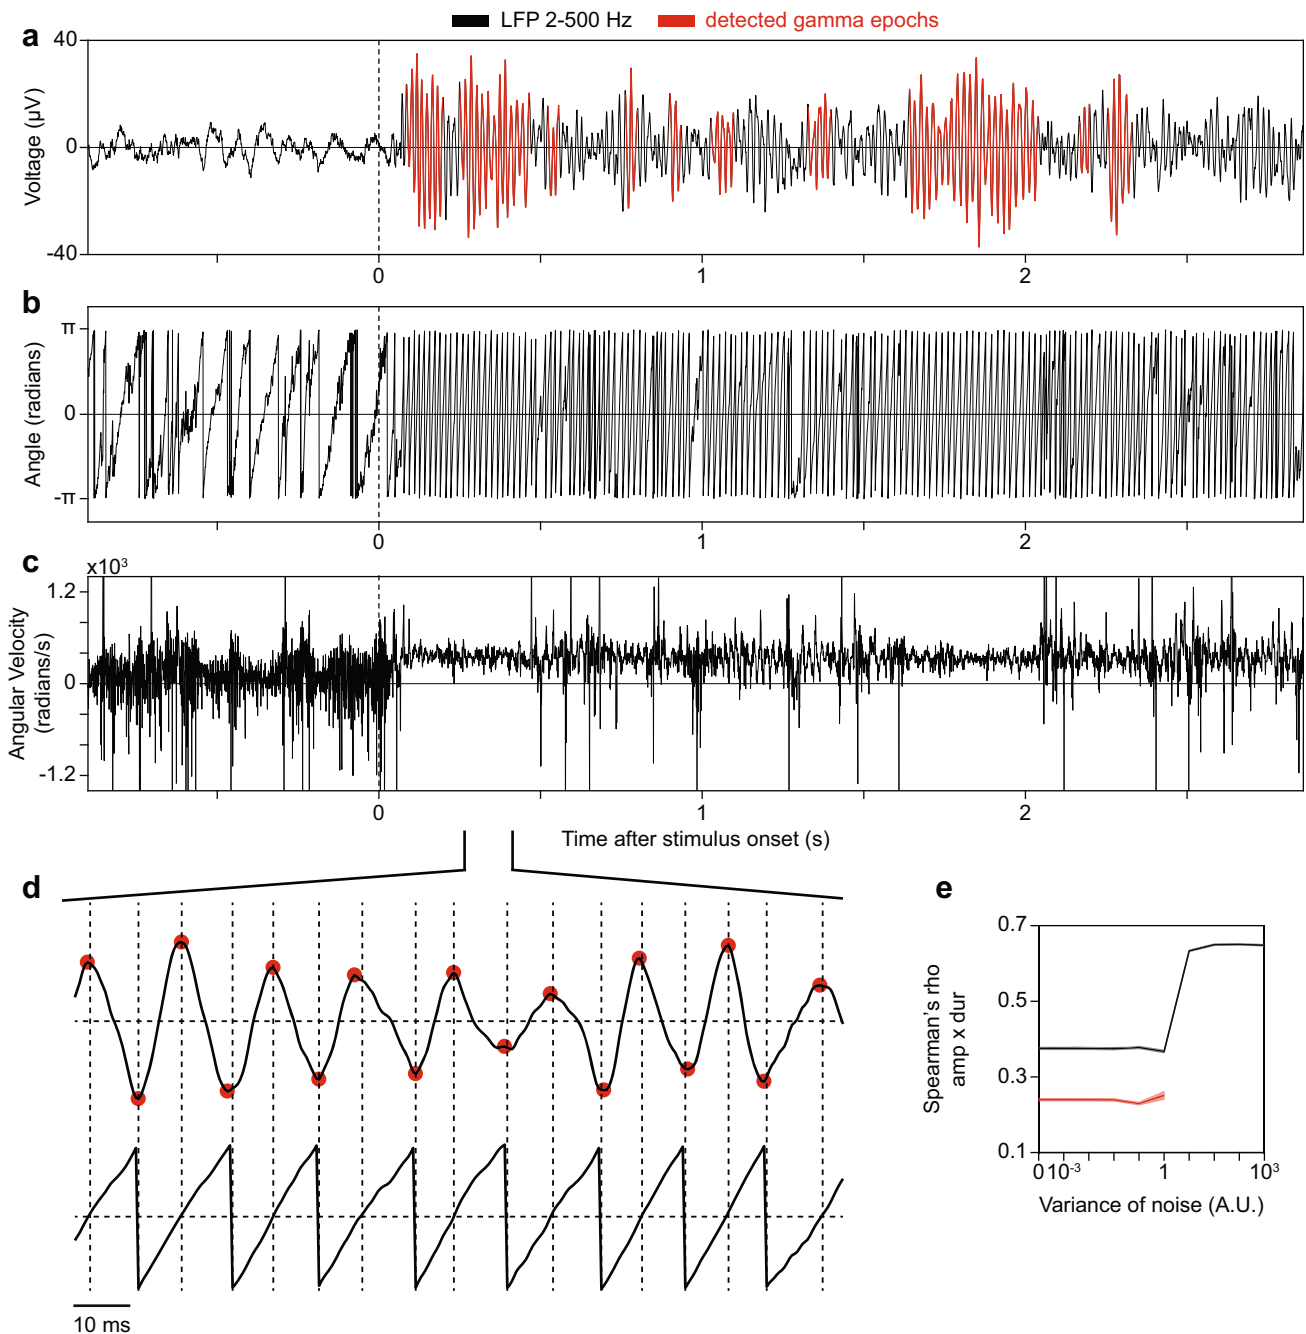
Fig. 3 Illustration of a method for the selection of gamma-oscillatory epochs. a LFP trace displayed in Fig. 1a, with regions presented in red corresponding to gamma epochs passing the criterion for stationarity. **b** Phase of the analytic signal based on the Hilbert transform of the trace shown in **a**. **c** Angular velocity of **a**. Note periods of positive and relatively stable angular velocity, corresponding to oscillatory gamma epochs in the original LFP. **a**–**c** Dashed lines indicate stimulus onset. **d** Magnification of the designated section of the LFP trace and its phase. Red dots indicate detected LFP peaks and troughs. Vertical dashed lines designate negative- to- positive and positive-to-negative zero crossings of the phase of the analytic signal, whereas horizontal dashed lines designate 0. **e** The correlation between the amplitude of a gamma half- cycle and the duration of the same gamma half cycle for different additive noise levels, computed with the method used in Fig. 2 (black), and the method described in this figure (red). The error regions show ±2 SEM based on a bootstrap procedure.

between the amplitude of a given half-cycle and the duration of multiple preceding and succeeding half-cycles (Fig. 4b). Some datasets showed dynamics on the temporal scale of few half-cycles (Fig. 4b, left), and others on the scale of multiple half-cycles (Fig. 4b, middle and right). For example, the right panel in Fig. 4b shows a long-lasting, negative trend punctuated by a small positive value for the instantaneous correlation. By contrast, the middle panel shows a positive trend peaking at zero lag. These

trends may have contributed to the observed correlation between gamma-cycle amplitude and duration. We therefore removed the influence of slower dynamics through a linear regression analysis (see Methods section). In this analysis, we first regressed out linear predictions of gamma-cycle amplitude and of duration from previous and succeeding cycles, obtaining regression residuals for amplitude and duration. We then repeated the analysis of amplitude-duration correlations on these regression residuals
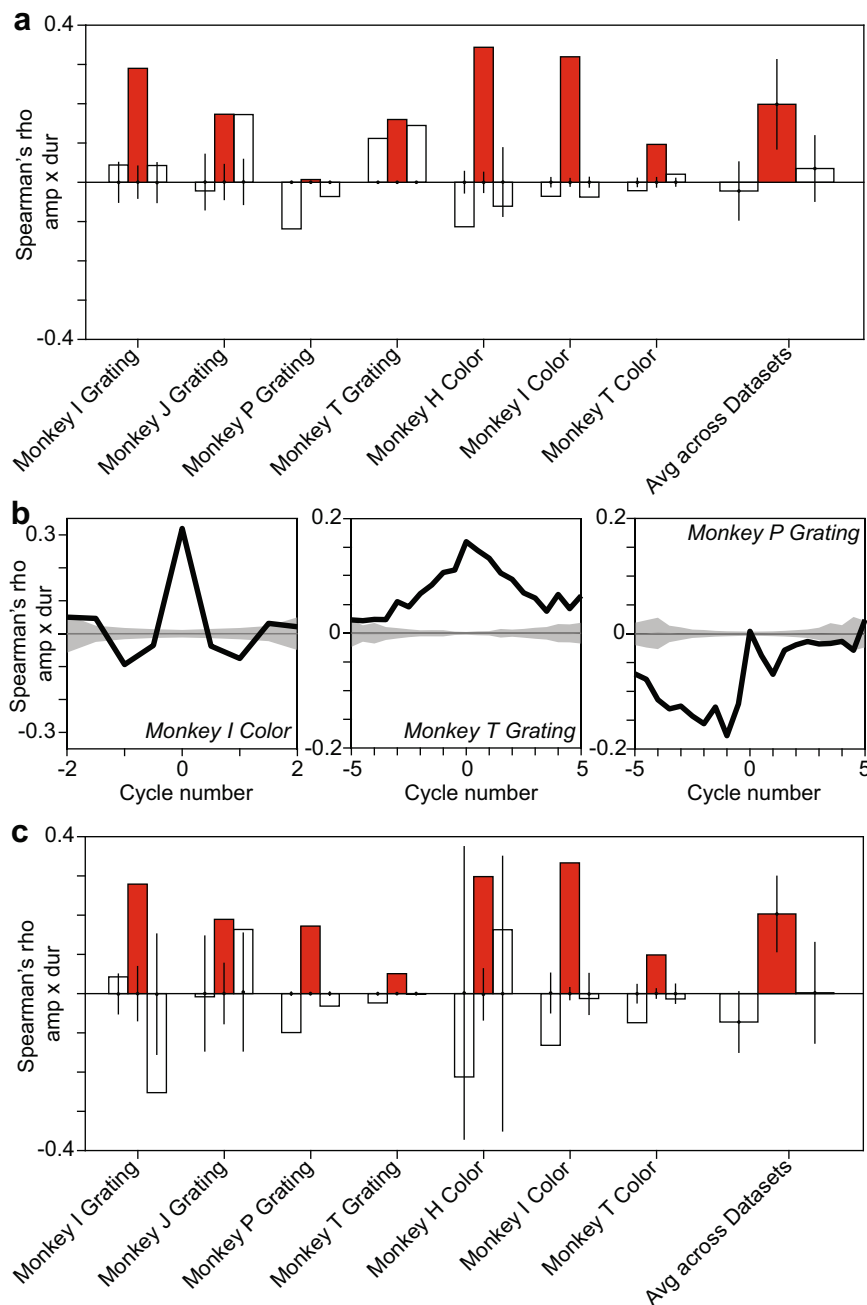
Fig. 4 Gamma-half-cycle amplitudes and durations are positively correlated in gamma-oscillatory epochs. a For each dataset listed on the x-axis, the three bars show the correlation between the amplitude of a gamma half cycle and the duration of the same (center, red), previous (left, white), and next (right, white) gamma half cycle. On the right, this is shown for the average across all datasets. This was calculated for each time point across trials and averaged across time points for gamma-oscillatory epochs. The data used correspond to the period during the presentation of the visual stimulus. For individual datasets, a null distribution was produced by randomizing the order of duration values across trials, and the resulting means and 99.9% confidence intervals are shown as dots and vertical lines (numbers of trials = 278, 5740, 672, 1075, 672, 320, 142 for respective datasets). For the average across datasets, shown on the right, we performed a two-sided *t*-test and show the resulting confidence intervals as vertical lines on the observed mean (red bar: $p < 6 \times 10^{-3}$, white bars for preceding cycle: $p = 0.5$, white bars for succeeding cycle: $p = 0.35$). In addition we performed a two-sided non-parametric permutation test (red bar: $p < 0.05$; white bars: $p > 0.05$). **b** Correlation between the amplitude of a gamma half-cycle and the duration of gamma half-cycles before and after it for three different datasets. Note that in monkey I, this is limited to ±2 cycles, because the signal-to-noise ratio was lower, resulting in shorter gamma-oscillatory epochs. Importantly, all three example datasets show a central peak, despite the fact that they show different longer-term correlations. The gray lines and gray-shaded areas depict the means and 99.9% confidence regions, after randomizing the order of duration values across trials. **c** Same as **a**, but showing the correlations between residuals of the regression across adjacent amplitude triplets and the residuals of the regression across adjacent duration triplets (numbers of trials = 278, 5740, 672, 1075, 672, 320, 142 for respective datasets; red bar: $p = 23 \times 10^{-4}$, two-sided t-test across datasets; $p < 0.05$, permutation test for individual datasets; white bars $p = 0.066$ and $p = 0.97$, respectively for preceding and succeeding cycles, two-sided *t*-test across datasets; $p > 0.05$, permutation test for individual datasets). **a–c**: $n = 6$ biologically independent experiments. Source data are provided in the Source Data file.
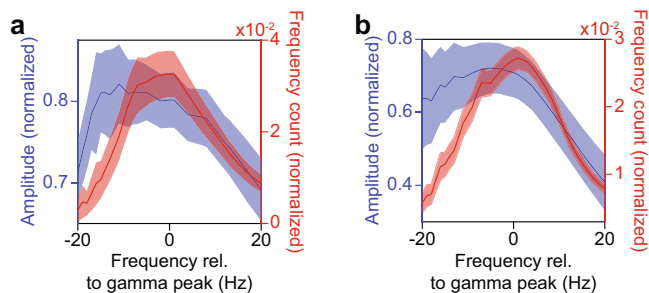
Fig. 5 Cycle-based amplitude-spectra and cycle-frequency distributions. a The x-axis shows duration expressed as its inverse, namely frequency, and after aligning to the gamma peak in the raw power spectrum. The blue curve shows the gamma-half-cycle amplitudes as a function of their duration. The red curve shows the count of detected gamma half-cycles as a function of their duration. These analyses were based on the broadband signal from the last 250 ms of stimulation (see Methods section). Error regions show ±2 SEM based on a bootstrap procedure. **b** Same as **a**, but for gamma epochs detected on the filtered LFP. Source data are provided in the Source Data file.

(see Methods section). We found that the resulting correlation was comparable to the correlation between raw amplitude and duration (compare Fig. 4a, c). Similar results were obtained for full cycles (Supplementary Fig. 1b). Together, these findings indicate that the positive correlation between gamma-cycle amplitudes and durations was not due to within- or across-trial trends on a longer timescale. Further analyses also suggest that the correlation between gamma-cycle amplitudes and durations was not due to transient changes in amplitudes and durations following microsaccades (Supplementary Fig. 2; see Methods section).

To further understand the contribution of non-stationarities to the correlation between gamma-cycle amplitudes and durations, we fitted autoregressive (AR) models to the LFP data (50–100 linear terms; see Methods section). AR models capture the variance and auto-correlation of the LFP, and can then be used to generate a stationary surrogate time-series, (by stationary we mean that the underlying statistics of the signal do not change over time). Supplementary Fig. 3a–d illustrates this for the dataset used for Fig. 1a–e. We find that the AR model accurately captured the power spectrum (Supplementary Fig. 3b), but did not replicate slower dynamics in gamma-cycle amplitudes or durations (Supplementary Fig. 3c, d; compare to Fig. 1d, e). In the surrogate data generated by the fitted AR models, we then analyzed the correlations between gamma-cycle amplitudes and durations, and found consistently positive correlations of similar average strength as in the original data (Supplementary Fig. 3e). Similar results were obtained for full rather than half cycles (Supplementary Fig. 1d). These results further support the notion that the observed correlations in the LFP data were not due to co-fluctuations or non-stationarities on a slower time scale.

The cycle-based amplitude spectrum and the rate of incidence of cycle-durations. In the V1 LFP data, we observed a small but positive correlation between gamma-cycle amplitudes and durations. This correlation, however, does not necessarily imply a monotonic or linear relationship between gamma-cycle amplitudes and durations, as was previously reported for hippocampus[24]. We thus examined the joint distribution of gamma-cycle amplitudes and durations in more detail. To this end, we first computed the average half-cycle amplitude for each possible half-cycle duration (Fig. 5a, see Methods section); we refer to this as the cycle-based amplitude spectrum (CBAS). To

minimize the possible influence of stimulus-locked trends in gamma amplitude and frequency, we used only the final 250 ms of visual stimulation. To average CBASs across monkeys, we first converted half-cycle duration values to frequency values (in Hz). We then aligned the CBASs to the "gamma peak frequency", that is the frequency at which the Fourier-based power spectrum (FBPS) reached a maximum. In the CBAS, we found that the relationship between frequency and amplitude was non-monotonic: The amplitude was greatest at a frequency that was lower than the peak gamma frequency, and showed a decline towards higher gamma frequencies. We further wondered how often different gamma-cycle durations tended to occur. We therefore computed the cycle- frequency (i.e. inverse of gamma-cycle duration) distribution. We found that the cycle-frequency distribution was approximately symmetric and was unimodal, which argues against multiple sources of gamma being mixed in. Specifically, the most prevalent half-cycle frequency lied within one Hertz of the peak gamma- frequency derived from the FBPS (Fig. 5a and Supplementary Fig. 4a).

We wondered whether the observed dependency of gamma-cycle amplitude on cycle-frequency may have been due to a ceiling effect, because our new method selected those broad-band LFP segments for which gamma rhythms were relatively strong. This selection circumvented several methodological problems, but it may have limited the generalizability of our findings. To address this issue, we re-analyzed the data after band-pass filtering the LFP in the gamma-frequency range (20–100 Hz). This modification substantially increased our sensitivity in detecting gamma episodes. The distributions of cycle-frequency and amplitude that we obtained after band-pass filtering were, nevertheless, highly similar to the ones calculated on the broad-band signal (Fig. 5b and Supplementary Fig. 4b).

Relationship of gamma frequency with spiking. Next, we aimed to obtain further insight into the mechanisms underlying amplitude-duration correlations. We first asked how the spontaneous dynamics of gamma oscillations related to the activation and phase-locking of excitatory and inhibitory neurons. The model of Atallah and Scanziani, discussed above, predicts that higher-amplitude gamma cycles are initiated by a stronger bout of excitatory spiking. These excitatory bouts should then give rise to longer-lasting inhibition, resulting in longer gamma cycles[24].

In order to assess if this prediction holds true for V1, we analyzed multi-unit (MUA) activity (see Methods section). We first computed the normalized spike count (number of spikes per cycle) (Fig. 6a) as a function of the gamma-cycle frequency (the inverse of gamma-cycle duration; see Methods section). The normalized spike count was negatively correlated with gamma-cycle frequency (Fig. 6d). This may be a trivial result, because the spike count reflects the product of firing rate and gamma-cycle duration. To correct for this, we divided the spike count by the duration, yielding the firing rate (spikes/sec). Firing rates were positively correlated with gamma-cycle frequency (Fig. 6b, d). To investigate if the same result held true for different excitatory and inhibitory cell classes, we classified single units into three classes that were previously identified by Onorato et al. in the same dataset: NW-Burst, NW-Nonburst (NW: Narrow waveform) and BW (Broad Waveform) units. Previous studies suggest that NW-Burst and BW units correspond to putative pyramidal cells, whereas NW-Nonburst neurons correspond to putative fast-spiking interneurons[29,36,37]. We found that firing rates were positively correlated with frequency for all three classes, similar to the MUA (Fig. 6d).

The model of Atallah and Scanziani posits that higher-amplitude gamma cycles are initiated by a stronger bout of
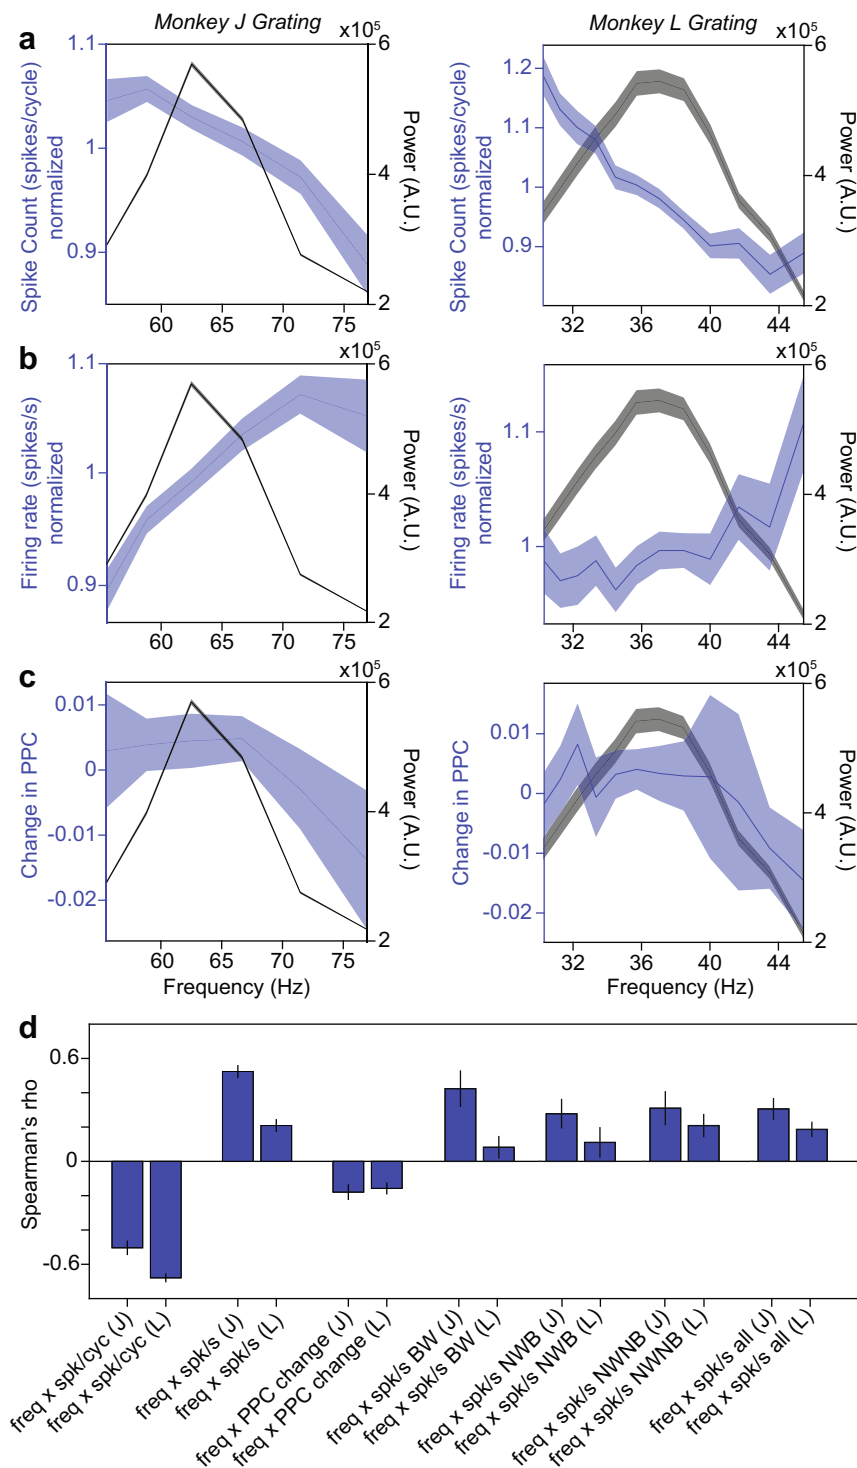
Fig. 6 The relationship between gamma-cycle duration and spiking. a The blue curve depicts the average normalized multi-unit (MU) spike count in detected gamma cycles of different durations, expressed on the x-axis as frequencies, for monkey J (left) and monkey L (right). The black curve depicts raw power in the gamma range of the respective monkeys. Error regions show ±2 SEM across units ($n = 156$ and 127 independent MU channels examined over 35 and 40 independent sessions in monkey J and L, respectively). **b** Same as **a**, but using the normalized MU firing rate. **c** Same as **a**, but showing the normalized change in spike-LFP PPC. **d** Correlation between the gamma-cycle duration, expressed as frequency, and several spiking metrics, separately for the two monkeys (J and L). Vertical lines depict ±2 SEM across units (no significant difference between cell-types, Kruskall–Wallis test, $p = 0.78$. SU: $n = 23$, 17, 16 independent SUs for NWB, NWNB, BW in monkey J, from 35 independent sessions; $n = 11$, 11, 13 NWB, NWNB, and BW in monkey L across 40 independent sessions, respectively. MU: same n's as in Fig. 6a). Source data are provided in the Source Data file.

excitatory spiking. Thus, it is possible that higher firing rates in the period of the LFP-trough to LFP-peak predict a longer duration of the next half cycle. We examined this possibility by computing the firing rate for the trough-to-peak period and correlating it with the duration of the next peak-to-trough half-cycle. Firing rate correlations were comparable to the same-cycle correlations and remained strongly positive for one monkey and non-significant for the other monkey (Supplementary Fig. 5).
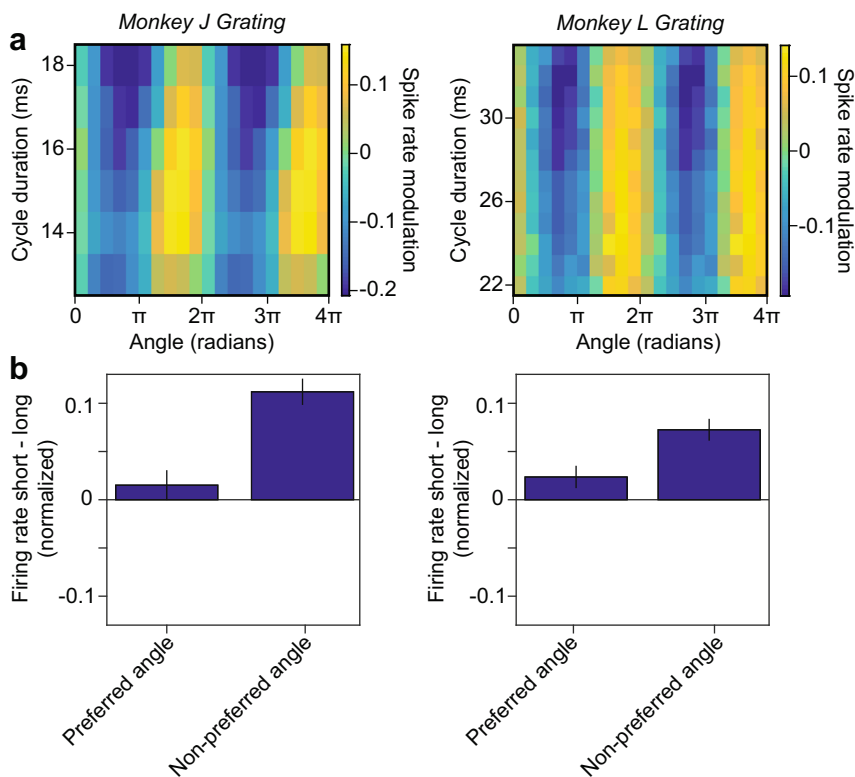
Fig. 7 The modulation of spiking activity by the phase of the gamma cycle. a The colormap shows the modulation of the MU firing rate as a function of gamma-cycle duration (y-axis) and the phase in the gamma cycle, at which spikes occurred (x-axis). Same n's as in Fig. 6a. **b** Difference in normalized firing rate between short and long gamma cycles for the preferred (left bar) and non-preferred phase in gamma cycles (right bar). Vertical lines depict ±2 SEM across units. Data from monkey J and monkey L are shown in the left and right column, respectively. Same n's as in Fig. 6a. The firing rate increase for short compared to long cycles at the non-preferred gamma phase is higher than at the preferred gamma phase both for monkey J and L (two-sided *t*-test: $p = 4.2 \times 10^{-6}$ and $p = 0.0029$ respectively). Source data are provided in the Source Data file.

These results are not in agreement with the predictions of the Atallah and Scanziani model.

We further wondered how spike synchrony was related to gamma-cycle duration. To investigate this, we (1) computed the duration of each gamma cycle, (2) identified all cycles of a certain duration, (3) pooled all spikes that were fired in those cycles together, and (4) computed spike-LFP phase-locking for each pool of spikes. We quantified phase locking with the pairwise phase consistency (PPC1)[38] measure, which removes potential biases due to spike count or firing history effects. We found that spike-LFP phase locking was negatively correlated with gamma frequency (Fig. 6c, d). The stronger spike-LFP phase locking in longer gamma cycles may have been due to a stronger spiking transient (at the "preferred" gamma phase), despite lower average firing-rates. To examine this, we divided each gamma cycle into eight non-overlapping phase-bins and computed MUA firing rates for these different bins. We did this separately for gamma cycles of different durations. As expected, longer gamma cycles showed a stronger phase modulation of firing rates (Fig. 7). However, we did not observe a stronger spiking transient in longer gamma cycles. Instead, in longer cycles, there was a stronger suppression of firing at the "non-preferred" gamma phase.

Gamma modeled by a damped harmonic oscillator driven by stochastic noise. We observed that correlations between gamma-cycle amplitudes and durations in awake macaque V1 are substantially weaker than reported by Atallah and Scanziani. Furthermore, the relation of instantaneous firing rate with gamma cycle duration, in particular the lack of a strong excitatory bout at

the onset of larger gamma cycles, is inconsistent with the model of Atallah and Scanziani. We thus wondered if a different model could explain our observations.

As a starting point for developing such a model, we used our observation (Supplementary Fig. 3) that positive correlations between gamma-cycle amplitudes and durations were also present for signals generated by the stationary AR model that we fitted to the LFP data (AR containing 50–100 linear terms). This observation was surprising for two reasons: (1) In an AR model, all variability in amplitude and duration is due to stochastic fluctuations in the innovation term (white noise); (2) In the AR model, all the interaction terms are linear (i.e. x[t] is a linear function of past values of x[t] plus white noise). To generate oscillatory behavior in an AR model, the minimum number of parameters that is required is two (AR(2)) (Supplementary Note 1). Further below we will show (see Supplementary Note 1) that AR(2) models can be interpreted as linear E-I circuits driven by stochastic drive. The characteristic behavior of an AR(2) can be described by its eigenvalues. For complex eigenvalues, the AR(2) model corresponds to a linear, damped harmonic oscillator that is stochastically driven (forced) by white noise. The strength of the oscillation is controlled by the magnitude of the eigenvalue (which needs to lie within the unit circle for the system to be stable) (Fig. 8a and see Supplementary Note 1). To directly compare AR(2) models to the LFP gamma oscillations, we fitted AR(2) models to the LFP power spectrum (see Methods). We found that AR(2) model fits could accurately reproduce the LFP power spectra in the gamma-frequency range (Fig. 8b), with eigenvalue magnitudes between approximately 0.97 and 0.995 (Fig. 8c), indicating that V1 gamma oscillations
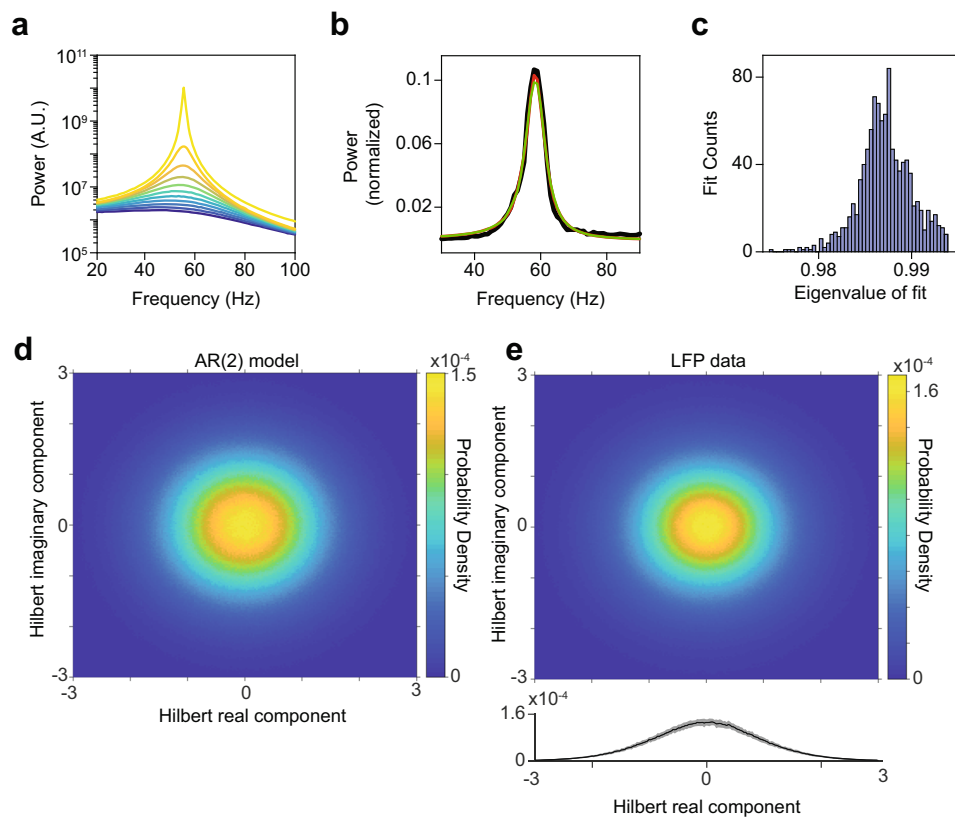
Fig. 8 AR(2) models fit to LFP data. a Power spectra of synthetic signals generated from AR(2) processes with corresponding eigenvalue magnitudes ranging from 0.9 (dark blue) to 0.999 (bright yellow). Note that we used a periodogram with a rectangular taper, in order to minimize the spectral leakage around the peak; this can introduce an amount of broad-band leakage. **b** Black: the change in LFP power relative to baseline as a function of frequency (Hz), for an example site in monkey T during the presentation of a full-screen drifting grating. Red: power spectrum of a synthetic signal generated by an AR(2) model. The AR(2) model was fitted to the LFP spectrum shown in black (see Methods section). Green: power spectrum of the I component of a linear PING model which is equivalent to the AR(2) model. **c** Histogram of eigenvalue magnitudes corresponding to AR(2) model fits of the LFP data. **d** Phase-plane of AR(2) data with eigenvalue magnitude of 0.99. Data were Hilbert-transformed. Shown are the imaginary and real component of the Hilbert-transform and their relative density. The imaginary and real component were z-scored to allow for averaging across channels. **e** Same as **d** but for LFP data. The LFP data were high-pass filtered above 10 Hz before being Hilbert-transformed. Shown below is the standard deviation across channels. Source data are provided in the Source Data file.

were close to criticality (note that AR(2) are unstable with eigenvalue magnitudes above 1). The phase-plane of an AR(2) model, obtained by examining the real and imaginary component of the analytical signal obtained via the Hilbert transform, exhibits a focus at the center of the phase plane with stochastic drive perturbing the equilibrium and causing damped oscillations. LFP data exhibited a similar phase-plane structure with a focus in the center (Fig. 8d, e).

Given the similarities between the spectral features of LFP and AR(2) data, we wondered whether the AR(2) model would produce positive correlations between gamma-cycle amplitudes and durations that are similar to the LFP data (note that the AR models investigated in Supp. Fig. 3 had 50–100 linear terms). We generated time series based on the AR(2) model and applied our method to detect gamma half-cycle amplitudes and durations. In these synthetic AR(2) signals, we found positive correlations between amplitudes and durations (Fig. 9a). For the same range of eigenvalue magnitudes, these correlations were comparable to the ones found in the V1 LFP data. Hence, positive correlations between gamma-cycle amplitudes and durations can be reproduced by a linear AR(2) model. Like the LFP data, AR(2) models also showed a non-monotonic relationship between gamma-cycle amplitude and duration and a roughly symmetric cycle-frequency distribution (Supplementary Fig. 6b). The standard deviation in gamma-cycle durations in the AR(2) model (~12 Hz for an

eigenvalue magnitude of 0.987, which is the median across LFP channels and stimulus conditions) matched well with the standard deviation in gamma-cycle durations for the LFP data (Supplementary Fig. 7; estimates ranged from 10 to 12.5 Hz, see Methods section).

The AR(2) model makes several additional predictions:

(1) It predicts that correlations between gamma-cycle amplitudes and durations should be smaller for conditions or channels in which gamma oscillations are on average stronger (Fig. 9a). We tested this prediction as follows: We first fitted an AR(2) model to the LFP spectra separately for each channel and condition, and determined the eigenvalue of the AR(2) fits (see Methods section). For the same LFP data, we then computed the amplitude-duration correlations, similar to Fig. 4. We then regressed the amplitude-duration correlation onto the eigenvalue magnitudes of the AR(2) fits (Fig. 9b). We found that, as predicted, amplitude-duration correlations decrease as a function of the eigenvalue magnitude. Together, these findings indicate that a simple AR(2) model predicts the observed amplitude-duration correlation and its negative dependence on average oscillation strength, as well as the joint amplitude-duration distribution.

(2) It predicts that amplitudes should be highly correlated across gamma cycles, i.e. there should be a very high autocorrelation of the gamma-cycle amplitude. This amplitude
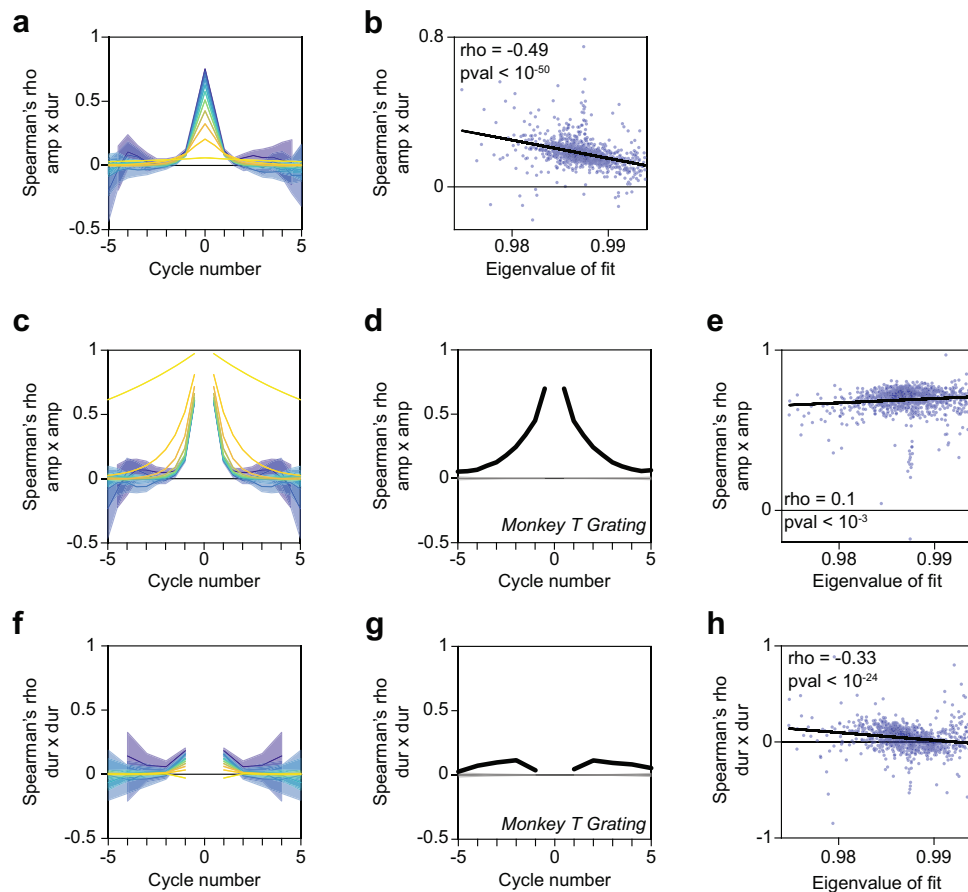
Fig. 9 A damped harmonic oscillator reproduces gamma-cycle amplitude and duration cross- and auto-correlations in the LFP. a Correlation between the amplitude of a gamma half-cycle and the duration of 10 gamma half-cycles before and after it. These correlation coefficients were computed for synthetic signals generated from AR(2) processes with corresponding eigenvalue magnitudes ranging from 0.9 (dark blue) to 0.999 (bright yellow) in steps of approximately 0.01. The error regions show ±2 SEM based on a bootstrap procedure. **b** Scatter plot of the eigenvalue magnitudes displayed in Fig. 8c and the instantaneous correlation between gamma half-cycle amplitude and duration from the corresponding LFP data. The regression fit (black line) was computed with the least-squares method. **c** The correlation between the amplitude of a given gamma half-cycle and the 10 gamma half-cycles before and after it (i.e. the autocorrelation) for synthetic signals generated from AR(2) processes. The error regions show ±2 SEM based on a bootstrap procedure. **d** The autocorrelation for LFP data from monkey T during presentation of a full-screen drifting grating. The gray lines and gray-shaded areas depict the means and 99.9% confidence regions, after randomizing the order of duration values across trials. **e** Same as **b**, but now showing the correlation between the amplitude of a given gamma half- cycle and the amplitude of its preceding and succeeding half-cycle, pooling data points from multiple datasets, conditions, and channels. **f**–**h** Same as **c**–**f** but for gamma full-cycle durations. Source data are provided in the Source Data file.

autocorrelation should be higher when gamma oscillations are on average stronger (Fig. 9c). We determined the amplitude autocorrelation by detecting the amplitude of all detected half-cycles in the LFP data. We then computed the autocorrelation between the amplitude of a given half-cycle with the amplitude of the previous and succeeding half-cycles. We found very high autocorrelations in half-cycle amplitude that were comparable to the ones observed in the AR(2) time series (Fig. 9d). Furthermore, as predicted, the amplitude autocorrelation was an increasing function of the eigenvalue magnitude (Fig. 9e). To rule out that the high amplitude correlations resulted from using half-cycles, we repeated this analysis on full cycles, and found essentially the same result (Supplementary Fig. 8).

(3) It predicts that gamma-cycle durations should be weakly correlated across gamma cycles, especially for strong gamma oscillations (Fig. 9f and Supplementary Fig. 9a). We computed autocorrelations based on the duration of all detected half-cycles in the LFP data. As predicted, the autocorrelation of the half-cycle durations was a decreasing function of the eigenvalue magnitude (Supplementary Fig. 9c). Yet, the autocorrelations of half-cycle durations were consistently negative (Supplementary Fig. 9b, c),

different from the autocorrelations in the AR(2) time series (which were positive or close to zero). This feature was likely due to asymmetric wave shapes of gamma cycles, which may perhaps reflect a difference in the time constants of the AMPA and GABA currents that generate the LFP and contribute to different parts of the gamma cycle. To avoid the potential influence of cycle asymmetry, we therefore repeated our analysis for full-cycle durations. We found that, as predicted, the autocorrelations of the full-cycle duration were positive but close to zero (bootstrap mean = 0.041; bootstrap SEM = 0.0038) (Fig. 9f–h) and that the autocorrelations were a decreasing function of the oscillation strength. In essence, this means that for strong oscillations, the deviation of the duration of the current gamma cycle from the mean gamma-cycle duration is not predictive of the deviation of the next gamma cycle.

Interpretation of AR(2) model and amplitude-duration correlations. Our results demonstrate that cycle-by-cycle dynamics of gamma oscillations are well reproduced by AR(2) models with complex eigenvalues. It remains unclear why neuronal interactions

between excitatory and inhibitory neurons in the local circuit would produce dynamics that resemble damped harmonic oscillators driven by noise. In the Supplementary Note 1, we demonstrate that AR(2) models can be written as a linear E-I circuit, by rewriting them as a system of two first-order stochastic difference equations with two populations (called "I" and "E") driven by external noise (which could e.g. correspond to stochastic inputs from the LGN)

$$I[t] = v_{ie}E[t-1] + v_{ei}I[t-1] + \varepsilon[t]$$
$$E[t] = v_{ee}E[t-1] + v_{ei}I[t-1] + \varepsilon[t]$$

(1)

If the AR(2) eigenvalues are complex conjugates, then the interaction terms $v_{ei}\, v_{ie} < 0$, i.e. the two state-variables (named I and E) always have opposite interaction weights (see Supplementary Note 1). Furthermore, the E population will always exhibit a phase lead over I. Given certain choices of the weight matrix **V**, this linear E-I circuit can reproduce a balanced E-I regime and a typical delay of E over I by a few ms[29,30,39,40]. In the Supplementary Note 1, we show how E and I can be estimated as a linear combination of the signal (x[t], x[t − 1]) through a transformation matrix **A** which corresponds to a weight matrix **V** (Supplementary Fig. 10, Supplementary Note 1 Supplementary Fig. 14). For example, if we let E[t] = x[t] − 1.1 x[t − 1] and I[t] = x[t] − 0.86x[t − 1], then the LFP has, as expected, opposite weights from E (−3.4) and I(+4.4), and the system exhibits E-I balance with a small phase-lead of E over I of 2.3 ms. In Supplementary Note 1 (Supplementary Fig. 13), we illustrate the behavior of the E and the I population for an AR(2) model with comparable eigenvalues to our data. We also simulated neuronal spikes based on inhomogeneous Poisson processes modulated by the E-variable in the AR(2) model and found that this simulation recapitulated the dependence of spike synchrony on cycle-duration (Supplementary Fig. 11).

Discussion

Circuits of excitatory and inhibitory neurons can generate rhythmic activity in the gamma frequency-range (30–80 Hz). Individual gamma-cycles show ample spontaneous variability in amplitude and duration. The mechanisms underlying this variability are not fully understood. We recorded local-field-potentials (LFPs) and spikes from awake macaque V1, and developed a noise-robust method to detect gamma-cycle amplitude and duration. This method circumvents several problems that arise due to band-pass filtering and peak/trough detection and allowed us to analyze the precise way in which amplitude and duration vary between gamma cycles, and how this variation relates to neuronal spiking activity. Our analyses establish several properties of gamma-oscillatory dynamics in macaque V1:

(1) The amplitude and duration of individual gamma cycles showed a weak but positive correlation (Spearman's rho = 0.199).

(2) Correlations between amplitude and duration were weaker for channels and conditions in which gamma oscillations were on average stronger.

(3) Gamma-cycle amplitude was strongly positively autocorrelated across cycles, especially for gamma oscillations that were on average stronger.

(4) Gamma-cycle duration was very weakly autocorrelated across gamma cycles, especially for gamma oscillations that were on average stronger. Thus, the deviation of the duration of the current gamma cycle from the mean gamma-cycle duration is not predictive of the deviation of the next gamma cycle.

(5) Longer gamma cycles were associated with stronger spike-field phase-locking (synchrony), but lower firing-rates, and were not accompanied by stronger, transient spiking activation.

We showed that the first four properties can be reproduced by random fluctuations in a system with resonance: A damped harmonic oscillator driven by stochastic noise (AR(2) model with complex roots). This model can be accurately fitted to V1 LFP data and is equivalent to a basic, linear E-I circuit driven by stochastic noise. Note that the idea that brain rhythms can be well modeled as damped harmonic oscillators was introduced many decades ago by Walter Freeman and others[41,42].

Atallah and Scanziani previously reported a strong positive correlation ($r = 0.61$) between gamma-cycle amplitude and duration in rat hippocampus. Here, we demonstrate that these positive correlations can arise due to the employed analysis method and the presence of noisy fluctuations in the signal (Fig. 2d, g). To avoid this problem, we developed an algorithm for the detection of gamma-oscillatory epochs, i.e. periods in the LFP dominated by gamma oscillations. Correlations computed for these periods remained positive, but were substantially weaker (Spearman's rho = 0.199; comparable result for Pearson's r) compared to Atallah and Scanziani[24]. Our analyses revealed several problems in the detection of gamma-cycle amplitude and frequency, due to the presence of non-stationarities in the LFP, and filter-generated smearing between adjacent data points in the time domain. This does not mean that our method detects the "ground-truth" gamma-cycle amplitude or duration: these quantities do not describe statistical properties that can be estimated, in contrast to quantities like the power spectral density. In a damped harmonic oscillator driven by noise, the notion of a "cycle" becomes fuzzy for low amplitudes: fluctuations become noise-driven, and the Hilbert-transform can yield negative frequencies, i.e. phase slips. For this reason, our cycle-detection method explicitly rejects epochs with phase slips as in[43].

There are several points of debate about the mechanisms of gamma oscillations. First, it remains unclear in which circuits, and under which conditions, gamma oscillations can be generated, and whether they primarily rely on I-I interactions (ING: Interneuron Network Gamma) or E-I interactions (PING)[3,6,25–29]. Several studies have observed a delay between E and I neurons (or intracellular E/I currents), consistent with the PING mechanism[29,39,40,44,45]. However, not all studies find such a phase delay[22,46,47]. Moreover, both PING and ING models can produce a wide range of dynamics depending on the specific parameter settings[28]. Second, the relative contributions of Somatostatin-positive (SSt +) and Parvalbumin-positive (PV +) interneurons remain unclear[6,48]: Whereas most theoretical and experimental studies support a general role of PV + interneurons in generating gamma oscillations, some circuits like mouse V1 exhibit oscillations around 30 Hz which may depend on SSt+ interneurons[48]. However, gamma oscillations in primate V1 have peak frequencies up to 70–80 Hz, and it is unlikely that these are mediated by SSt+ interneurons given the relatively long membrane time constants of SSt+ interneurons and the lack of temporally precise SSt+ responses to inputs from principal cells[49]. Third, in primate and cat V1, there exist specialized excitatory neurons that do not exist in mouse V1 and that may play an important role in generating high-amplitude gamma oscillations[30,36].

Here, we show that many features of gamma-oscillatory dynamics in awake macaque V1 are predicted from a stationary model containing only linear dynamics. In this model, oscillations emerge from perturbations by stochastic drive. It is often assumed that variability in gamma-cycle amplitude and duration results from non-linear dynamics or non-stationarities in the underlying signal, e.g. due to eye movements[21] or cross-frequency coupling[50]. However, we show that spontaneous variability in gamma-cycle amplitude and duration in monkey V1 is consistent with a stationary AR(2) model, which is equivalent to a linear E-I circuit driven by stochastic inputs. To produce amplitude-duration correlations, this model does not require the presence of a strong, transient bout of excitatory activity to produce long gamma cycles, as was supposed by the non-linear E-I model of

Atallah and Scanziani, but which was not observed in our data (Fig. 7). Rather, our analyses indicate that correlations between amplitude and duration result from filtering random fluctuations by a damped harmonic oscillator. Given that the expected velocity in the phase-plane of the AR(2) model is constant and independent of radius, correlations between amplitude and duration arise from the interaction between stochastic drive and the linear deterministic dynamics of the system.

We have shown that gamma dynamics in macaque V1 are well reproduced by damped harmonic oscillators driven by noise, which are equivalent to linear E-I circuits. Our model thereby connects two lines of research on gamma dynamics: On the one hand models in which gamma results from E-I interactions[51,52]. On the other hand, the model of gamma as filtered white noise[4], which, like the AR(2) model, is also a stationary signal model that reproduces the power spectrum of the signal. (Note that while the AR(2) is a form of filtered white noise, the reverse is not necessarily the case). Burns et al. showed that the distribution of gamma-burst durations can be reproduced by generating filtered white-noise, i.e. a mix of sinusoids with random phases and the same amplitude as the LFP power spectrum (which is different from an AR(p) model)[4]. Further, by computing auto-coherence over the wavelet-transform amplitudes of the LFP signal, Burns et al.[53] found an auto-coherence of gamma over time with a resultant length of 0.3–0.4, similar to auto-coherence values observed in filtered white noise. Here we performed a related analysis with a cycle-by-cycle detection method that avoids spurious correlations due to windowing or band-pass filtering. We found that the correlation between the current and the next cycle-duration is close to zero (bootstrap mean = 0.041), and approaches zero for strong oscillations. We find that the standard deviation in gamma-cycle frequency is around 10 Hz even for very strong oscillations (with eigenvalue magnitudes ~0.99) (note that this only holds for the time periods exhibiting clear cycles).

The stochastic nature of gamma oscillations may have implications for their putative role in inter-areal communication (refs. [54–57]; see Supplementary Discussion), neural representation and working memory. Our findings on stochasticity in the amplitude and duration of individual gamma cycles can inform the debate regarding whether neural representations are persistent or transient in nature. There is evidence that fluctuations in gamma-amplitude, frequency and phase are accompanied by fluctuations in, neuronal tuning, neuronal correlations and behavioral performance. Shortened reaction times in a given trial can be predicted by enhanced gamma power in human visual cortex[58], enhanced gamma synchronization in macaque V4[59] and enhanced gamma coherence between macaque V1 and V4[60]. Further evidence shows that inter-areal information transmission fluctuates with the relative gamma phase[60–62]. In terms of local information processing, both orientation selectivity and noise correlations of V1 firing rates fluctuate as a function of gamma phase and amplitude[63]. These stochastic fluctuations in sensory information processing and transmission might indicate the need for relatively long integration times for stable coding and communication performance[5], or alternatively might reflect that information processing occurs not in a persistent but in a rhythmic manner[13]. The stochastic nature of gamma also features in a recent debate on the persistent versus transient nature of working memory representations. Recent evidence shows that working memory representations become transiently activated during short gamma bursts[11]. A possible reconciliation between the persistent/transient perspectives might be that there is persistent encoding of working memory information in the underlying statistics of the population signals, but that working memory information is predominantly transmitted

during oscillatory bursts which can naturally emerge due to stationary, stochastic fluctuations. Interestingly, damped harmonic oscillators exhibit memory in the sense that the past sequence of noise inputs is stored in the ongoing amplitude of the signal. Hence, when oscillations are strong, they would "store" potential energy that can be released in the form of a damped oscillation, like a mass on a spring. During these oscillatory bursts, the system might be primarily governed by deterministic dynamics releasing energy, whereas during low-amplitude fluctuations, the system trajectories might predominantly follow external inputs.

It remains unclear whether our simple model of a damped harmonic oscillator driven by noise reproduces all features of gamma-oscillatory dynamics; it is possible that more complex models are needed in order to do so, and our model primarily models spontaneous fluctuations in gamma dynamics. However, it is quite surprising that several aspects of the gamma oscillations in the collective, high-dimensional dynamics of millions of V1 neurons, measured at the macro/meso-scale, are well predicted from a model that is linear and contains only two parameters.

Methods

Subjects. We analyzed data from a total of 6 adult macaque monkeys (*macaca mulatta*), referred to as monkey H, I, J, L, P, and T. Monkeys I and L are/were female, the others male. The experiments were approved by the responsible regional or local authority, which was the Regierungspräsidium Darmstadt, Germany, for monkeys H, I, J, L, and T, and the ethics committee of the Radboud University, Nijmegen, Netherlands, for monkey P.

Recordings. We used different recording procedures and stimulus paradigms for the different monkeys, and will describe these separately for the different monkeys.

Task. All monkeys performed a passive fixation task. The specific details of the task performed by monkeys I and P were as follows: Monkeys initiated a trial by depressing a lever (monkey I) or touching a bar (monkey P), which triggered the appearance of a fixation point, and then brought their gaze into a fixation window around the fixation point. Monkeys were required to fixate on the fixation point, which was centered on a gray background, after which a stimulus was presented. If they kept their gaze within the fixation window as long as the stimulus was presented, they were given a juice reward after the release of the lever/bar following stimulus offset. Monkeys H, J, L, and T performed a similar task, with the initiation/termination of the trial being solely dependent on the acquisition/release of fixation (i.e. not dependent on pressing a lever or touching a bar). Further details of this version of the task are described in Peter et al.[33] for monkey H, and in Lima et al.[19] for monkeys J and L. For all monkeys, fixation windows ranged between 0.5 and 1.2 degrees radius.

Recordings (electrodes, reference). For monkey H, recordings were done with CerePort ("Utah") arrays (64 micro-electrodes; inter-electrode distance 400 μm, tip radius 3–5 μm, impedances 70–800 kΩ, half of them with a length of 1 mm and half with a length of 0.6 mm, Blackrock Microsystems). A reference wire was inserted under the dura toward parietal cortex. Further details are reported in[33]. For monkey I, a semi-chronic microelectrode array micro-drive was implanted over area V1 of the left hemisphere (SC32-1 drive from Gray Matter Research; 32 independently movable glass insulated tungsten electrodes with an impedance range of 0.5–2 MΩ and an inter-electrode distance of 1.5 mm, electrodes from Alpha Omega). We used the micro-drive chamber as the recording reference. For monkeys J and L, recordings were performed with 2–10 microelectrodes, made of quartz-insulated, tungsten-platinum material (diameter: 80 μm; impedances between 0.3 and 1 MΩ; wire from Thomas Recording). These were inserted independently into the cortex via transdural guide tubes (diameter: 300 μm; Ehrhardt Söhne), which were assembled in a customized recording device (designed by S.N.). This device consisted of 5 precision hydraulic micro-drives mounted on an X-Y stage (MO-95, Narishige Scientific Instrument Laboratory, Japan), which was secured on the recording chamber by means of a screw mount adapter. Inter-electrode distance ranged between 1 and 3 mm. We used the micro-drive chamber as the recording reference. Further details are reported in Lima et al.[19]. For monkey P, we recorded neuronal activity with a micro-machined 252-channel electro-corticogram (ECoG) electrode array implanted subdurally on the left hemisphere[64,65]. We used a silver ball implanted over occipital cortex of the right hemisphere as the recording reference. For monkey T, we recorded neuronal activity with a micro-machined 252-channel ECoG electrode array implanted subdurally over areas V1 and V4 of the left hemisphere (inter-electrode distance 1400 μm; electrode diameter 400 μm, IMTEK & BCF, University of Freiburg)[65]. We

used an electrode adjacent to the lunate sulcus as a recording reference for the section of the array covering area V1.

Recordings (acquisition, filtering). For monkeys H, I, and T, we acquired data with Tucker Davis Technologies (TDT) systems. Data were filtered between 0.35 and 7500 Hz (3 dB filter cutoffs) and digitized at 24,414.0625 Hz (TDT PZ2 pre-amplifier). For monkeys J and L, we obtained spiking activity and the LFP by amplifying 1000 times and band-pass filtering (0.7–6.0 kHz for MUA; 0.7–170 Hz for LFP) with a customized 32-channel Plexon pre-amplifier connected to an HST16o25 headstage (Plexon Inc., USA). Additional 103-fold signal amplification was performed by onboard amplifiers (E-series acquisition boards, National Instruments, USA). For monkey P, we acquired data with a Neuralynx system. Data were amplified 20 times, high-pass filtered at 0.159 Hz, low-pass filtered at 8 kHz, and digitized at 32 kHz by a Neuralynx Digital Lynx system.

Receptive field mapping/eccentricities. Receptive fields (RFs) were mapped with either bar stimuli (refs. [19,33]; monkeys H, I, J, L), patches of moving gratings (refs. [64]; monkey P) or red dots (monkey T). The signal used for RF mapping was multi-unit activity (MUA) for monkeys H, I, J, L, and the LFP gamma power for monkeys P and T. For monkeys J and L, we recorded neuronal activity from the opercular region of area V1, leading to RF-center eccentricities of 2–3 deg, and occasionally from the superior bank of the calcarine sulcus, leading to RF-center eccentricities of 10-13 deg. For monkey H, RF-center eccentricities ranged between 5.2 and 7.1 deg (median RF-center eccentricity 6.2 deg). For monkey I, RF-center eccentricities ranged between 2.6 and 6.7 deg (median RF-center eccentricity 4.5 deg). For monkey P, RF-center eccentricities ranged between 3 and 5.7 deg (median RF-center eccentricity 4.6 deg). For monkey T, RF-center eccentricities ranged between 3.1 and 7.1 deg (median RF-center eccentricity 3.8 deg).

Eye position monitoring. For monkeys H, I and T, eye movements and pupil size were recorded at 1000 Hz using an Eyelink 1000 system (SR Research Ltd.) with infrared illumination. For monkeys J and L, we monitored the eye position with a scleral search coil system (DNI, Crist Instruments, USA; sampling rate of 500 Hz). For monkey P we monitored eye position with an infrared camera system (Thomas Recording ET-49B system) at a sampling rate of 230 Hz. We used a standardized fixation task in order to calibrate eye signals before each recording session.

Behavioral control and stimulus presentation. Stimulus presentation and behavioral control was implemented as follows: The software toolbox ARCADE (https://gitlab.com/esi-neuroscience/arcade) was used for monkeys H, I, and T; Custom LabVIEW code (Lab-VIEW, National Instruments, USA) was used for monkeys J and L; The software toolbox CORTEX (dally.nimh.nih.gov/index.html) was used for monkey P.

Monkeys H and I were presented with full-screen uniform color surfaces. Surface color varied across trials according to a pseudo-random sequence. For our analyses, we used the hue that elicited the strongest gamma oscillations (monkey H RGB: 149 99 0; monkey I RGB: 255 0 0). In a separate session, monkey I was also repeatedly presented with a full-screen drifting square-wave red-and-green grating of a fixed initial phase and drift-direction (RGB for red 255 0 0 and green 0 255 0; spatial frequency: 1.5 cycles/degree; temporal frequency 2 Hz). Monkeys J and L were presented with large drifting square-wave black-and-white gratings (spatial frequencies: 1.25–2 cycles/degree; temporal frequencies: 1.4–2 Hz) and plaid stimuli. Only the gratings were used for our analyses. The gratings had a diameter of 8 degrees of visual angle and were positioned at the average of the RF centers of the recorded MUA. In each trial, the direction of the grating drift was randomly chosen from 16 directions (in steps of 22.5 degrees). Monkey P was repeatedly presented with a full-screen drifting square-wave black-and-white grating of a fixed initial phase and drift-direction (spatial frequency: ~1 cycle/degree; temporal frequency ~1 Hz). Monkey T was presented with full-screen uniform color surfaces, with the color changing across trials according to a pseudo-random sequence. For our analyses, we used two hues that elicited the strongest gamma oscillations (RGB: 255 0 0 and 0 0 255). In separate sessions, monkey T was also presented with full-screen drifting square-wave colored gratings of pseudo-random initial phases and drift-directions. For our analyses, we used the gratings that elicited the strongest gamma oscillations (red-green RGB: 255 0 0 and 0 255 0 and blue-yellow RGB: 0 0 255 and 255 255 0; spatial frequency: 1.5 cycles/degree; temporal frequency 2 Hz). For monkeys H, I and T, stimuli were presented on 120 Hz LCD monitors[66], without gamma correction. For monkeys J, L and P, stimuli were presented on CRT monitors (100-120 Hz), after gamma correction.

Data analysis. All analyses were done in MATLAB (The MathWorks) using custom scripts and the FieldTrip toolbox (www.fieldtriptoolbox.org[67]). The analyses were done only on trials with correct task performance. In monkeys P and T, we selected the 25% electrodes/sites over area V1 with the strongest visually induced gamma band activity, because the grids covered a relatively large region of retinotopic space and contained electrodes that were poorly driven by the visual stimulus. In monkeys H, I, J and L, we analyzed all visually driven electrodes. In all monkeys except for monkey T, we analyzed LFP signals that were recorded relative to the common reference signal (described above). For monkey T, we calculated

local bipolar derivatives between LFPs from immediately neighboring electrodes. i.e., differences (sample-by-sample in the time domain), similar to previous studies[64]. This was done because the global references in monkey T were positioned over V1 and V4 in the same hemisphere.

Preprocessing. For monkeys H, I, and T, LFPs were obtained from the broadband signal after low-pass filtering (sixth order Butterworth filter with a corner frequency of 500 Hz), high-pass filtering (third order Butterworth filter with a corner frequency of 2 Hz for monkey T and 4 Hz for monkeys H and I) and down-sampling to 2034.51 Hz. For monkeys J and L, LFPs were filtered between 0.7 and 170 Hz (hardware-filter, described above) and down-sampled to 1 kHz. For monkey P, we obtained LFP signals by low-pass filtering at 200 Hz and down-sampling to 1 kHz. In addition, for monkey P, we removed powerline artifacts at 50 Hz and its harmonics with a digital notch filter.

Segmenting data into epochs, and calculation of power and TFR. To estimate the LFP power spectra in the stimulus and baseline periods (Figs. 1b, c, g, h, 5, and 6a–c, and Supplementary Figs. 3b and 4), we used the following procedure: Power spectra were estimated separately for the pre-stimulus period and the stimulation period. The pre-stimulus period was the time between fixation onset and stimulus onset. During the pre-stimulus period, monkeys fixated on a central dot on a gray screen, and there was no other stimulus presented. For monkeys H, I, P, and T, the pre-stimulus and stimulation periods were of variable length across trials. We kept data corresponding to the pre-stimulus and stimulation period with the minimum length (monkey H: baseline 0.3 s/stimulation 1.5 s; monkey I: baseline 0.5 s/stimulation 2 s; monkey P: baseline 0.3 s/stimulation 2.3 s; monkey T: baseline 1.1 s/stimulation with full-screen gratings 2.8 s/stimulation with full-screen uniform color surfaces 3.2 s). For monkeys J and L, the pre-stimulus and grating-stimulation periods had a stable duration across trials within a session but their duration varied between sessions. All of the available pre-stimulus and grating data were analyzed for those monkeys (baseline 0.8–1 s/stimulation 2–2.4 s). The power spectral analysis was based on epochs of fixed lengths. Therefore, the described task periods were cut into non-overlapping epochs. We aimed at excluding data soon after stimulus onset ("event") to minimize the influence of the stimulus-onset related event-related potential on our analyses. Therefore, periods were cut into non-overlapping epochs, starting from the end of the period and stopping before an epoch would have included data ~0.5 s after those events. For Fig. 1b, c, g, h, the estimation of power spectra was based on epochs of 0.5 s length; for Figs. 5 and 6a–c and Supplementary Figs. 3b and 5, power spectra were based on epochs of 0.25 s. Data epochs were Hann tapered, to achieve a fundamental spectral resolution (Rayleigh frequency) of 2 Hz (4 Hz for Figs. 5 and 6a–c and Supplementary Figs. 3b and 5), and then Fourier transformed. The gamma-band power spectra used for the AR(2) fits (Figs. 8b, c and 9b, e, h and Supplementary Figs. 8b and 9c), the power spectra of synthetic AR(2) signals (Fig. 8a), and the joint distribution of gamma-cycle amplitude and duration (Supplementary Fig. 6) were based on rectangular windows of 1 s, in order to ensure minimal spectral smearing, and thus a more accurate fit. For the time-frequency analysis of power, we used window lengths of ±2.5 cycles per frequency which were slid over the available data in steps of 1 ms. Power during the stimulation period was normalized to the pre-stimulus baseline period, separately for each channel, in the following manner: Power per frequency and per trial was calculated as described above. Power calculated for the pre-stimulus baseline period was then averaged across trials. Finally, trial-wise normalized power was calculated for the stimulation period by subtracting the average pre-stimulus spectrum and then dividing by it.

Spike sorting. Single units were isolated through semi-automated spike sorting[30]. First, we performed semi-automatic clustering with the KlustaKwik 3.0 software. The energy of the spike waveform and the energy of its first derivative were used as features in this procedure. A candidate single unit was accepted if the corresponding cluster was clearly separable from the noise clusters, and if the inter-spike-interval distribution had a clear refractory-period. This was done manually with the M-Clust software. In addition, we used the isolation distance (ID[68]) as a measure of cluster separation. The ID of a candidate single unit had to exceed 20 in order for it to be included in our analyses. The median ID was 25.05. This procedure led to the isolation of 100 single units. For each isolated single unit, we computed the peak-to-trough duration of the average AP waveform. Single units with long (>0.235 ms) and short (<0.235 ms) peak-to-trough durations were named "broad-waveform" (BW) and "narrow-waveform" (NW) neurons, respectively. Broad-waveform neurons corresponded to 29% of the single unit population.

Initial estimation of gamma-cycle amplitude and duration (cf. Atallah and Scanziani). For our initial analyses of individual gamma cycles, we implemented the algorithm as described by Atallah and Scanziani for data from awake freely-moving rats. In short, we first low-pass filtered the LFP by using a 40 ms moving average filter and then subtracted this filtered signal from the original time series (Experimental Procedures and Supplemental Experimental Procedures of Atallah and Scanziani), which effectively corresponds to a high-pass filter with a corner frequency at ~20 Hz. The resulting signal was further band-pass filtered in the range of 5–100 Hz with a 3rd order, two-way Butterworth filter. Gamma-cycle

peaks and troughs were then defined as local maxima and minima, respectively. Furthermore, gamma-cycle amplitudes were defined as the difference between the voltage of a given peak and its subsequent trough. Similarly, gamma-cycle durations were defined as the interval between a given peak and it subsequent peak. This analysis was done in segments of the filtered signal which displayed high power in the individual gamma frequency range of each dataset (peak gamma frequency ±20 Hz). These segments were extracted in the following way: a time-power representation of each trial was calculated with 5 discrete prolate slepian sequences and windows of 100 ms which were slid over the available data in steps of 25 ms. Gamma episodes were defined as segments of the resulting time-series which lasted for more than 100 ms and had power that exceeded a threshold. This threshold was calculated separately for each trial as the difference between the mean of the time-power representation and its standard deviation.

Generation of colored noise. In Fig. 2g, we analyzed the correlations obtained with the Atallah-Scanziani method for colored noise. We generated noise with power spectra following a $1/f^n$ function, where f denotes frequency and n assumes 11 equally spaced values between, and including, 0 (corresponding to white noise) and 2 (corresponding to Brownian noise). This was done in the following manner: (i) 1000 white noise traces containing 10^6 samples were generated for each n. (ii) Each trace was Fourier transformed. (iii) The complex coefficients of the positive frequencies in the resulting spectra were multiplied by the $1/f^n$ function. (iv) A synthetic spectrum was constructed by concatenating the above complex coefficients with the conjugate of their flipped version. (v) The resulting spectrum was inverse Fourier transformed to obtain time series.

Improved estimation of gamma-cycle amplitude and duration. We developed an improved method to extract gamma-cycle amplitude and frequency from the LFP signals as follows:

(1) We computed the Hilbert-transform of the broadband LFP signal to obtain the analytic signal and derive the time-resolved phase from it. We used the broadband signal, because band-pass filtering creates dependencies between voltage values across time points, and can transform transient, non-oscillatory deflections into rhythmic events.

(2) We detected gamma cycles as follows: First, we detected all the zero-crossings of the phase. Such phase zero crossings occur in the neighborhood of peaks and troughs in the original LFP signal. For each k-th zero-crossing, we examined whether the angular velocity of the phase was positive for all time points between the k − 1-th to the k + 1-th zero-crossing (similar to Muller et al.[69]). If this was not the case, then there was a negative "phase-slip" in which the instantaneous frequency became negative, and the respective zero crossing plus/minus two neighboring zero crossings were discarded. Negative instantaneous frequencies make the interpretation of the instantaneous frequency and amplitude ambiguous, and are typically accompanied by small peaks/troughs in the LFP signal. This violates our model of the gamma oscillation as a signal with a positive frequency which fluctuates over time, $y(t) = A(t) * \cos(\omega(t)*t + \varphi)$, where $A(t)$ and $\omega(t)$ are the instantaneous amplitude and frequency fluctuating over time.

If there was no negative phase-slip, then we identified gamma peaks by first detecting negative-to-positive zero crossings in the phase of the analytic signal. For each of these crossings, we then identified the nearest local maximum in the LFP signal (Fig. 3d). Likewise, gamma troughs were identified by detecting positive-to-negative zero crossings and identifying nearby local minima. Using the detected gamma peaks and troughs, we then determined the gamma-cycle amplitude and duration. To obtain estimates of gamma-cycle amplitude and duration with the maximum attainable temporal resolution, we divided each gamma cycle into "half-cycles": The first half-cycle comprised the data segment from the trough to the peak, and the second half-cycle from the peak to the trough. For each half-cycle, amplitude was defined as the difference between the respective peak and trough, and duration was defined as the corresponding time interval. For each detected half-cycle, we thus obtained an amplitude and duration value. For comparison, we also determined amplitude and duration for full gamma cycles. A gamma cycle comprised the data from one peak to the next peak. Amplitude was defined as the voltage difference between the first peak and the trough. Duration was defined at the time between the two peaks.

Note that for the analysis of the relationship between individual gamma cycles and spiking activity, we used a band-pass filter (3rd order, two-pass Butterworth, with a pass-band of 40–90 Hz for monkey J and 25–55 Hz for monkey L). In this case, we used an additional criterion to reject epochs of spurious rhythmic activity[30]: We ran the same cycle-selection procedure on the pre-stimulus period, in which narrow-band gamma-band oscillations are virtually absent. For the pre-stimulus period, we obtained the mean μ_{pre} and standard deviation σ_{pre} of the distribution of amplitudes. These amplitudes were measured as the peak-to-trough distance of the gamma cycle. A cycle in the stimulus period with amplitude A was only selected if $(A - \mu_{pre})/\sigma_{pre} > 1.63$ (which is equivalent to a one-sided T-test at $P < 0.05$). We filtered the LFP with the purpose of increasing the number of selected gamma epochs, considering that the analysis of unit firing rates and spike-field phase-locking demands a relatively large amount of data. Note that we have shown in Fig. 5 that the distributions of amplitude and frequency after band-pass filtering are comparable to the distributions obtained without band-pass filtering. In addition, the potential issues related to filtering only apply to the calculation of

correlations of amplitude and duration and not to the calculation of the correlation of spiking strength and gamma frequency. This is due to the fact that filtering may generate artificial correlations between the amplitudes and durations of deflections of the same time series (explained further in the results section). The filter used on the LFP is not used on the spiking activity. Thus, artificial correlations between spiking and cycle-by-cycle frequency are not likely.

Amplitude and frequency values were extracted from selected gamma epochs of a duration of at least two full cycles.

Computation of time-resolved correlations between amplitude and frequency. In the case of our V1 recordings, we observed that gamma amplitude and cycle duration progressively increased over time after the onset of a drifting grating stimulus. (Fig. 1c, d). By contrast, after the onset of a uniform color surface, gamma amplitude and duration progressively decreased and increased over time, respectively (Fig. 1g, h). These changes with time after stimulus onset could contribute to the correlation values between gamma-cycle amplitude and duration, if gamma amplitude and duration values are concatenated across all trials and time points. This would conceal the relationship between gamma-cycle amplitude and duration due to intrinsic variability, by introducing a positive or negative correlation bias for drifting gratings and uniform color surfaces, respectively.

We avoided these effects by using the following method: We calculated correlations between gamma-cycle amplitudes and durations across all trials, separately for each time point (at the respective sampling rate) after stimulus onset, and subsequently averaged those correlation values over time points and subsequently over recording sites. To enable this, we needed to define gamma-cycle amplitudes and durations for each time point. Therefore, each time point (relative to stimulus onset) was localized to the gamma half cycle (or full cycle), into which it fell, and it was assigned the respective amplitude and duration of that half cycle (or full cycle). For the calculation of correlations with one or multiple half-cycle (or full-cycle) lags, correlations were calculated between amplitudes and durations shifted relative to each other by the corresponding number of half-cycles (or full cycles).

In datasets containing more than one stimulus condition, correlation coefficients were calculated separately for each condition and then averaged across conditions.

As mentioned in the results section, the correlation analysis used the Spearman correlation coefficient. Like in Atallah and Scanziani[24], we found results to be essentially identical for Spearman and Pearson correlation, when using their method of determining gamma-cycle amplitude and duration. For the rest of our analyses, we used exclusively the Spearman correlation coefficient.

Statistical significance of correlations. The statistical significance of auto- and cross- correlations of gamma-cycle amplitudes and durations, and correlations between AR(2)-fit eigenvalue-magnitudes and auto- or cross correlations of gamma-cycle amplitudes and durations was assessed by means of a non-parametric randomization approach. In this paragraph, we will describe this approach for the cross-correlation of amplitudes and durations: The order of valid duration values was randomly shuffled across trials, separately for each time-point. We then computed surrogate Spearman's correlation coefficients 1000 times as described above for each dataset. Next, we performed a fit of a Gaussian distribution on the 1000 surrogate correlation coefficients. Empirical correlations were deemed significant if they were 3 standard deviations larger or smaller than the mean of the surrogate distribution. This procedure implements a non-parametric version of a two-sided test with a p-value of ≈ 0.001.

To test if the mean correlation of gamma-cycle amplitudes and durations is significantly different from zero across datasets, we applied a Student's t-test. In general, we prefer non-parametric randomization tests over parametric tests (like the t-test). However, some analyses contained only four or five datasets, which effectively precludes the application of non-parametric tests. Where possible, we supplemented the t-test with a non-parametric statistical test (Figs. 2c, 4a, c, and Supplementary Fig. 1a). Specifically, we calculated the mean correlation across datasets for each possible combination of values that results after independently inverting or maintaining the sign of each correlation value (i.e. a full permutation). This led to a surrogate distribution of mean values to which the empirical mean was compared for statistical significance. Mean correlations were deemed significant if they were larger (smaller) than the top (bottom) 2.5 percentile of this surrogate distribution.

Regression analysis. We performed regression analyses separately for gamma-cycle amplitudes and durations with the Matlab function *regress*. As explained in the results section, for each half-cycle, we regressed the amplitude value of the ongoing half-cycle against the amplitude values of the previous and next half-cycle, by using a least squares approach. We used the same procedure for half-cycle duration values. This was done for each time point after stimulus onset separately, and by using all the amplitude and duration values across trials (for that time point).

For Fig. 4c, we first calculated the regression residuals for amplitude and duration separately by predicting amplitude and duration from the previous and next half cycle. We then computed the regression residuals for amplitude and duration, separately for each time point. These residual values measured the extent to which the amplitude or duration in the ongoing half-cycle was greater or smaller

than predicted from the surrounding half-cycles, and thereby departed from slower trends. We then computed the correlation between the regression residuals for amplitude and duration, in the same way as described above.

Micro-saccade detection. We low-pass filtered vertical and horizontal eye position signals by replacing each value with the average over itself ±15 ms. We then computed the first temporal derivative of the signals to obtain the vertical and horizontal velocities. We combined those values to obtain the eye speed irrespective of the direction of eye movement. Per trial, we determined the SD of eye speed, and any deviation >4 SDs and lasting for at least 30 ms was deemed a saccadic eye movement. Saccadic eye movements that remained within the fixation window were considered to be MSs.

AR. In Supplementary Fig. 3, we computed our correlations for data generated through auto-regressive models with a power spectrum similar to the recorded LFP data. An autoregressive (AR) model of order n represents each value in a time-varying process as the linear sum of its n preceding values (each weighted by a separate coefficient) and a stochastic term. This model can then be used to generate a synthetic time series that has the same power spectrum as the original process, but that is devoid of higher-order statistical properties such as slow temporal trends or spectral cross-frequency dependencies. We modeled the LFP as an AR process of a relatively high order (50 for monkeys J and P, whose analysis was based on a sampling rate of 1000 Hz, and 100 for monkeys H, I, T, whose analysis was based on a sampling rate of 2034.51 Hz). We did this by fitting a vector of AR coefficients and a noise variance term with the Matlab function *arfit*, simultaneously to all the trials of a given stimulus condition and independently for each recording site. For our analyses, we only used the period of the trial starting at 250 ms after stimulus onset, thereby omitting stimulus onset-related transient activity. These AR models were then used to generate surrogate time series.

AR(2) Signal generation and model derivation. For the AR(2) model relationship to E-I circuits we refer to Supplementary Note 1.

To generate AR(2) signals, we computed the AR(2) coefficients for a given eigenvalue magnitude and oscillation frequency, using standard analytical transformations. Generated time series were analyzed with the same cycle-detection method as the LFP data. The only difference was that for the AR(2), we did not divide the data into trials, and thus computed the correlation between cycle amplitude and duration across all the cycles over all the time points (i.e. not across trials for each time point separately). In order to compare the AR models to the LFP data, we ensured that the model used a sampling frequency of 2035 Hz, similar to the sampling frequency of most of our LFP datasets. For several analyses, we correlated the eigenvalue-magnitude of the AR(2) fit to the LFP data, with several correlation measures across LFP datasets, including the amplitude-duration correlation, amplitude autocorrelation and duration autocorrelation. To ensure that all preprocessing (sampling rate; filtering) was similar for these data, we only included datasets with a similar sampling frequency of 2034.51 Hz.

AR(2) Model Fit to LFP data. We estimated the strength of gamma oscillations in our LFP data as follows: (1) We computed gamma-band power spectra separately for each channel and condition. The power spectra were based on rectangular windows of 1 s, in order to ensure minimal spectral smoothing, and thus a more accurate fit. (2) We then estimated the coefficients of equivalent AR(2) models by minimizing the squared error in the gamma frequency-range (*matlab* function *fminsearch*) between each LFP power spectrum and the following function:

$$S(f) = \frac{\sigma_z^2}{1 + \varphi_1^2 + \varphi_2^2 - 2\varphi_1(1 - \varphi_2)\cos(2\pi f) - 2\varphi_2\,\cos(4\pi f)}$$

where $S(f)$ is the power spectrum of the AR(2), σ_z is the standard deviation of this power spectrum, f are frequencies in the gamma range, and φ_1/φ_2 are the AR(2) coefficients (Fig. 8c). (3) We determined the eigenvalues of the equivalent AR(2) models (Figs. 8b, c and 9b, e, h, and Supplementary Figs. 8b and 9c).

PPC. For the calculation of spike-LFP PPC, the gamma phase of each spike within a gamma cycle was defined as $t/T*2*\pi$, where t was the time of the spike relative to the start of the gamma cycle, and T was the duration of the gamma cycle. This constitutes a linear phase interpolation. This used the improved Hilbert-based definition of gamma half-cycles (cycles). The obtained spike phases from separate trials were collected, and the average consistency of phases across these pairs was estimated with the pairwise-phase-consistency metric (PPC)[38,70], and more specifically its PPC1 variant[70]. Any potential bias due to differences in discharge rates is removed by the pairwise computation. Only neurons that fired at least 50 spikes were considered, because phase-locking estimates can have a high variance in cases of low spike counts. We were not able to perform this analysis for single-unit activity, due to the lack of a sufficient number of detected single unit spikes.

Computation of the cycle-based amplitude spectrum (CBAS) and cycle-frequency distribution. For Fig. 5 and Supplementary Fig. 4 we computed the cycle-based amplitude spectrum (CBAS) and the cycle-frequency distribution as

follows. Gamma half-cycle amplitude and duration values were extracted from the LFP through the use of the previously described improved detection algorithm. Values of gamma-half-cycle durations were converted into values of gamma-half-cycle frequency (frequency being the inverse of duration). This was done separately for each recording site and stimulus condition. Next, gamma half cycles were assigned to their corresponding frequency bin, and for each frequency bin, the average amplitude and the rate of incidence of that frequency were determined.

Note that the peak gamma-frequency varies across experimental subjects and stimulus conditions. In order to compute averages across stimulus conditions and monkeys, it is therefore necessary to align individual distributions to the power-spectral peak in the gamma-frequency-range, separately for each stimulus condition and dataset. We performed this alignment in the following way: The raw trial-wise power spectra were estimated separately for each stimulus condition as described above (see power), and from these spectra we determined the peak gamma-frequency. In addition, this was done for the baseline-corrected power spectra. The alignment of half-cycle amplitudes and frequency counts was then performed around the resulting frequency. Specifically, half-cycle amplitude and frequency count averages at ±20 Hz around the gamma peak were averaged across stimulus conditions and datasets. Note that we analyzed datasets with different sampling rates. This entailed that the range of detectable half-cycle frequencies (i.e. sampling rate/(2*duration)) varied across different datasets and, depending on the sampling rate, certain frequency bins were necessarily empty. In order to average across datasets with different sampling rates, we therefore performed a linear interpolation between normalized half-cycle amplitude values and frequency counts, which were adjacent to empty bins.

Estimation of gamma cycle-frequency variability. Gamma-cycle frequency variability was quantified as the standard deviation of the distribution of gamma-cycle frequencies. This standard deviation was computed in two alternative ways: first, we performed Gaussian fits on the cycle-frequency distributions of LFPs and an AR(2) model (eigenvalue magnitude of 0.987, which is the median across LFP channels and stimulus conditions). This approach yielded a standard deviation of 10.5017 and 11.7199 Hz for the LFP data, and 11.9724 Hz for the AR(2) model (Supplementary Fig. 7). Second, we computed the standard deviation of gamma cycle-frequencies in the LFP by using either individual pairs or triplets of adjacent gamma cycles. In the analysis of gamma-cycle pairs we implemented Bessel's correction, where the variance between the cycle-frequency of adjacent cycle-pairs is computed, this pairwise variance is averaged across cycle-pairs, and the final estimate of the standard deviation is derived by the square root of this average. The analysis of gamma-cycle triplets differed from the analysis of pairs only in that it involved estimating the pairwise variance of the middle cycle of each triplet to the mean of its two neighboring cycles (interpolation). Note that this analysis requires further debiasing by a scalar which can be analytically proven to be exactly 2/3. These analyses led to an estimate of 12.52 Hz and 12.3154 Hz for pairs and triplets, respectively, which is very similar to the estimates based on Gaussian fits to cycle-frequency distributions.

Reporting summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

Source Data is provided for all figures. The data sets generated in this study are under active use for future publications by the reporting laboratory. Accordingly the relevant data sets presented in this manuscript will be shared upon reasonable request. Inquiries for data access should be sent to the corresponding authors and will typically be responded to within a matter of days. Source data are provided with this paper.

Code availability

MATLAB code used in this study is available upon reasonable request. Source data are provided with this paper.

References

1. Singer, W. & Gray, C. M. Visual feature integration and the temporal correlation hypothesis. *Annu. Rev. Neurosci.* **18**, 555–586 (1995).
2. Roelfsema, P. R., Lamme, V. A. & Spekreijse, H. Synchrony and covariation of firing rates in the primary visual cortex during contour grouping. *Nat. Neurosci.* **7**, 982–991 (2004).
3. Tiesinga, P. & Sejnowski, T. J. Cortical enlightenment: are attentional gamma oscillations driven by ING or PING? *Neuron* **63**, 727–732 (2009).
4. Burns, S. P., Xing, D. & Shapley, R. M. Is gamma-band activity in the local field potential of V1 cortex a "clock" or filtered noise? *J. Neurosci.* **31**, 9658–9664 (2011).

ARTICLE

5. Akam, T. E. & Kullmann, D. M. Efficient "communication through coherence" requires oscillations structured to minimize interference between signals. *PLoS Comput. Biol.* **8**, e1002760 (2012).

6. Buzsáki, G. & Wang, X. J. Mechanisms of gamma oscillations. *Annu. Rev. Neurosci.* **35**, 203–225 (2012).

7. Schomburg, E. W. et al. Theta phase segregation of input-specific gamma patterns in entorhinal-hippocampal networks. *Neuron* **84**, 470–485 (2014).

8. Fries, P. Rhythms for cognition: communication through coherence. *Neuron* **88**, 220–235 (2015).

9. Hermes, D., Miller, K. J., Wandell, B. A. & Winawer, J. Stimulus dependence of gamma oscillations in human visual cortex. *Cereb. Cortex* **25**, 2951–2959 (2015).

10. Ray, S. & Maunsell, J. H. Do gamma oscillations play a role in cerebral cortex? *Trends Cogn. Sci.* **19**, 78–85 (2015).

11. Lundqvist, M. et al. Gamma and beta bursts underlie working memory. *Neuron* **90**, 152–164 (2016).

12. Vinck, M. & Bosman, C. A. More gamma more predictions: gamma-synchronization as a key mechanism for efficient integration of classical receptive field inputs with surround predictions. *Front. Syst. Neurosci.* **10**, 35 (2016).

13. Palmigiano, A., Geisel, T., Wolf, F. & Battaglia, D. Flexible information routing by transient synchrony. *Nat. Neurosci.* **20**, 1014–1022 (2017).

14. Cardin, J. A. Snapshots of the brain in action: local circuit operations through the lens of gamma oscillations. *J. Neurosci.* **36**, 10496–10504 (2016).

15. Sohal, V. S. How close are we to understanding what (if anything) gamma oscillations do in cortical circuits? *J. Neurosci.* **36**, 10489–10495 (2016).

16. Freeman W. J. *Mass Action in the Nervous System* (Academic Press, 1975).

17. Varela, F., Lachaux, J. P., Rodriguez, E. & Martinerie, J. The brainweb: phase synchronization and large-scale integration. *Nat. Rev. Neurosci.* **2**, 229–239 (2001).

18. Buzsáki G. *Rhythms of the Brain* (Oxford University Press, 2005).

19. Lima, B., Singer, W., Chen, N. H. & Neuenschwander, S. Synchronization dynamics in response to plaid stimuli in monkey V1. *Cereb. Cortex* **20**, 1556–1573 (2010).

20. Ray, S. & Maunsell, J. H. Differences in gamma frequencies across visual cortex restrict their possible use in computation. *Neuron* **67**, 885–896 (2010).

21. Lowet, E., Roberts, M. J., Bosman, C. A., Fries, P. & De Weerd, P. Areas V1 and V2 show microsaccade-related 3-4-Hz covariation in gamma power and frequency. *Eur. J. Neurosci.* **43**, 1286–1296 (2016).

22. Perrenoud, Q., Pennartz, C. M. & Gentet, L. J. Membrane potential dynamics of spontaneous and visually evoked gamma activity in V1 of awake mice. *PLoS Biol.* **14**, e1002383 (2016).

23. Lowet E. et al. A quantitative theory of gamma synchronization in macaque V1. *Elife* **6**, e26642 (2017).

24. Atallah, B. V. & Scanziani, M. Instantaneous modulation of gamma oscillation frequency by balancing excitation with inhibition. *Neuron* **62**, 566–577 (2009).

25. Traub, R. D., Whittington, M. A., Colling, S. B., Buzsáki, G. & Jefferys, J. G. Analysis of gamma rhythms in the rat hippocampus in vitro and in vivo. *J. Physiol.* **493**, 471–484 (1996).

26. Wehr, M. & Zador, A. M. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* **426**, 442–446 (2003).

27. Okun, M. & Lampl, I. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat. Neurosci.* **11**, 535–537 (2008).

28. Wang, X. J. Neurophysiological and computational principles of cortical rhythms in cognition. *Physiol. Rev.* **90**, 1195–1268 (2010).

29. Vinck, M., Womelsdorf, T., Buffalo, E. A., Desimone, R. & Fries, P. Attentional modulation of cell-class-specific gamma-band synchronization in awake monkey area v4. *Neuron* **80**, 1077–1089 (2013).

30. Onorato, I. et al. A distinct class of bursting neurons with strong gamma synchronization and stimulus selectivity in monkey V1. *Neuron* **105**, 180–197 e185 (2020).

31. Roberts, M. J. et al. Robust gamma coherence between macaque V1 and V2 by dynamic frequency matching. *Neuron* **78**, 523–536 (2013).

32. Shirhatti, V. & Ray, S. Long-wavelength (reddish) hues induce unusually large gamma oscillations in the primate primary visual cortex. *Proc. Natl Acad. Sci. USA* **115**, 4489–4494 (2018).

33. Peter, A. et al. Surface color and predictability determine contextual modulation of V1 firing and gamma oscillations. *eLife* **8**, e42101 (2019).

34. Brunet, N. M. et al. Stimulus repetition modulates gamma-band synchronization in primate visual cortex. *Proc. Natl Acad. Sci. USA* **111**, 3626–3631 (2014).

35. Vinck, M., Batista-Brito, R., Knoblich, U. & Cardin, J. A. Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron* **86**, 740–754 (2015).

36. Gray, C. M. & McCormick, D. A. Chattering cells: superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science* **274**, 109–113 (1996).

37. Mitchell, J. F., Sundberg, K. A. & Reynolds, J. H. Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* **55**, 131–141 (2007).

38. Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P. & Pennartz, C. M. The pairwise phase consistency: a bias-free measure of rhythmic neuronal synchronization. *Neuroimage* **51**, 112–122 (2010).

39. Csicsvari, J., Jamieson, B., Wise, K. D. & Buzsáki, G. Mechanisms of gamma oscillations in the hippocampus of the behaving rat. *Neuron* **37**, 311–322 (2003).

40. Hasenstaub, A. et al. Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron* **47**, 423–435 (2005).

41. Freeman, W. J. Linear approximation of prepyriform evoked potential in cats. *Exp. Neurol.* **5**, 477–499 (1962).

42. Gersch, W., Yonemoto, J. & Naitoh, P. Automatic classification of multivariate EEGs using an amount of information measure and the eigenvalues of parametric time series model features. *Comput. Biomed. Res.* **10**, 297–318 (1977).

43. Davis, Z. W., Muller, L., Martinez-Trujillo, J., Sejnowski, T. & Reynolds, J. H. Spontaneous travelling cortical waves gate perception in behaving primates. *Nature* **587**, 432–436 (2020).

44. Börgers, C. & Kopell, N. Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neural Comput.* **15**, 509–538 (2003).

45. Tiesinga, P. H., Fellous, J. M., Jose, J. V. & Sejnowski, T. J. Computational model of carbachol-induced delta, theta, and gamma oscillations in the hippocampus. *Hippocampus* **11**, 251–274 (2001).

46. Siegle, J. H., Pritchett, D. L. & Moore, C. I. Gamma-range synchronization of fast-spiking interneurons can enhance detection of tactile stimuli. *Nat. Neurosci.* **17**, 1371–1379 (2014).

47. Vinck, M. et al. Cell-type and state-dependent synchronization among rodent somatosensory, visual, perirhinal cortex, and hippocampus CA1. *Front Syst. Neurosci.* **9**, 187 (2015).

48. Veit, J., Hakim, R., Jadi, M. P., Sejnowski, T. J. & Adesnik, H. Cortical gamma band synchronization through somatostatin interneurons. *Nat. Neurosci.* **20**, 951–959 (2017).

49. Jouhanneau, J. S., Kremkow, J. & Poulet, J. F. A. Single synaptic inputs drive high-precision action potentials in parvalbumin expressing GABA-ergic cortical neurons in vivo. *Nat. Commun.* **9**, 1540 (2018).

50. Spyropoulos, G., Bosman, C. A. & Fries, P. A theta rhythm in macaque visual cortex and its attentional modulation. *Proc. Natl Acad. Sci. USA* **115**, E5614–E5623 (2018).

51. Powanwe, A. S. & Longtin, A. Determinants of brain rhythm burst statistics. *Sci. Rep.* **9**, 18335 (2019).

52. Wallace, E., Benayoun, M., van Drongelen, W. & Cowan, J. D. Emergent oscillations in networks of stochastic spiking neurons. *PLoS ONE* **6**, e14804 (2011).

53. Burns, S. P., Xing, D. & Shapley, R. M. Comparisons of the dynamics of local field potential and multiunit activity signals in macaque visual cortex. *J. Neurosci.* **30**, 13739–13749 (2010).

54. Grothe, I., Neitzel, S. D., Mandon, S. & Kreiter, A. K. Switching neuronal inputs by differential modulations of gamma-band phase-coherence. *J. Neurosci.* **32**, 16172–16180 (2012).

55. Schneider M., Dann B., Sheshadri S., Scherberger H., Vinck M. A general theory of coherence between brain areas. *bioRxiv* https://doi.org/10.1101/2020.06.17.156190 (2020).

56. Pesaran, B. et al. Investigating large-scale brain dynamics using field potential recordings: analysis and interpretation. *Nat. Neurosci.* **21**, 903–919 (2018).

57. Buzsaki, G. & Schomburg, E. W. What does gamma coherence tell us about inter-regional neural communication? *Nat. Neurosci.* **18**, 484–489 (2015).

58. Hoogenboom, N., Schoffelen, J. M., Oostenveld, R. & Fries, P. Visually induced gamma-band activity predicts speed of change detection in humans. *Neuroimage* **51**, 1162–1167 (2010).

59. Womelsdorf, T., Fries, P., Mitra, P. P. & Desimone, R. Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* **439**, 733–736 (2006).

60. Rohenkohl, G., Bosman, C. A. & Fries, P. Gamma synchronization between V1 and V4 improves behavioral performance. *Neuron* **100**, 953–963 e953 (2018).

61. Womelsdorf, T. et al. Modulation of neuronal interactions through neuronal synchronization. *Science* **316**, 1609–1612 (2007).

62. Ni, J. et al. Gamma-rhythmic gain modulation. *Neuron* **92**, 240–251 (2016).

63. Womelsdorf, T. et al. Orientation selectivity and noise correlation in awake monkey area V1 are modulated by the gamma cycle. *Proc. Natl Acad. Sci. USA* **109**, 4302–4307 (2012).

64. Bosman, C. A. et al. Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron* **75**, 875–888 (2012).

65. Rubehn, B., Bosman, C., Oostenveld, R., Fries, P. & Stieglitz, T. A MEMS-based flexible multichannel ECoG-electrode array. *J. Neural Eng.* **6**, 036003 (2009).

66. Wang, P. & Nikolić, D. An LCD monitor with sufficiently precise timing for research in vision. *Front. Hum. Neurosci.* **5**, 85 (2011).
67. Oostenveld, R., Fries, P., Maris, E. & Schoffelen, J. M. FieldTrip: open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Comput. Intell. Neurosci.* **2011**, 156869 (2011).
68. Schmitzer-Torbert, N., Jackson, J., Henze, D., Harris, K. & Redish, A. D. Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience* **131**, 1–11 (2005).
69. Muller, L., Reynaud, A., Chavane, F. & Destexhe, A. The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave. *Nat. Commun.* **5**, 3675 (2014).
70. Vinck, M., Battaglia, F. P., Womelsdorf, T. & Pennartz, C. Improved measures of phase-coupling between spikes and the Local Field Potential. *J. Comput Neurosci.* **33**, 53–75 (2012).

Author contributions
Conceptualization by G.S., P.F., and M.V.; Methodology by G.S., M.S., P.F., and M.V.; Software, G.S., M.S., J.R.D., and M.V.; Analysis of LFP data, G.S. and M.V.; Analysis of spiking data, M.V.; Spike sorting, I.O. and M.V.; Simulations and mathematical analysis of AR(2) model, M.S, G.S., and M.V.; Recordings, G.S., J.R.D., M.L.S., C.A.B., B.L., A.P., J.K.-L., R.R., S.N., W.S., and P.F.; Writing, G.S., M.S., P.F., and M.V.; Supervision, P.F. and M.V.; M.S. and J.R.D. contributed equally to this study.

Competing interests
P.F. has a patent on thin-film electrodes and is beneficiary of a respective license contract with Blackrock Microsystems LLC (Salt Lake City, UT), he is member of the Scientific Technical Advisory Board of CorTec GmbH (Freiburg, Germany) and managing director of Brain Science GmbH (Frankfurt am Main, Germany). The authors declare no other competing interests.

Additional information
Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41467-022-29674-x.

Correspondence and requests for materials should be addressed to Georgios Spyropoulos, Pascal Fries or Martin Vinck.

Peer review information *Nature Communications* thanks Richard Naud, Vikaas Sohal and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permission information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.