The Magnitude, But Not the Sign, of MT Single-Trial Spike-Time Correlations Predicts Motion Detection Performance

Alireza Hashemi,¹ Ashkan Golzar,¹ Jackson E.T. Smith,² and Erik P. Cook¹

¹Department of Physiology, Centre for Applied Mathematics in Bioscience and Medicine, McGill University, Montreal, Quebec H3G 1Y6, Canada and

²Department of Physiology, Anatomy and Genetics, University of Oxford, Oxford OX1 3PT, United Kingdom

Spike-time correlations capture the short timescale covariance between the activity of neurons on a single trial. These correlations can significantly vary in magnitude and sign from trial to trial, and have been proposed to contribute to information encoding in visual cortex. While monkeys performed a motion-pulse detection task, we examined the behavioral impact of both the magnitude and sign of single-trial spike-time correlations between two nonoverlapping pools of middle temporal (MT) neurons. We applied three single-trial measures of spike-time correlation between our multiunit MT spike trains (Pearson's, absolute value of Pearson's, and mutual information), and examined the degree to which they predicted a subject's performance on a trial-by-trial basis. We found that on each trial, positive and negative spike-time correlations were almost equally likely, and, once the correlational sign was accounted for, all three measures were similarly predictive of behavior. Importantly, just before the behaviorally relevant motion pulse occurred, single-trial spike-time correlations were as predictive of the performance of the animal as single-trial firing rates. While firing rates were positively associated with behavioral outcomes, the presence of either strong positive or negative correlations had a detrimental effect on behavior. These correlations occurred on short timescales, and the strongest positive and negative correlations modulated behavioral performance by ~9%, compared with trials with no correlations. We suggest a model where spike-time correlations are associated with a common noise source for the two MT pools, which in turn decreases the signal-to-noise ratio of the integrated signals that drive motion detection.

Key words: area MT; electrophysiology; monkey; motion detection; neural correlations; visual cortex

Significance Statement

Previous work has shown that spike-time correlations occurring on short timescales can affect the encoding of visual inputs. Although spike-time correlations significantly vary in both magnitude and sign across trials, their impact on trial-by-trial behavior is not fully understood. Using neural recordings from area MT (middle temporal) in monkeys performing a motion-detection task using a brief stimulus, we found that both positive and negative spike-time correlations predicted behavioral responses as well as firing rate on a trial-by-trial basis. We propose that strong positive and negative spike-time correlations decreased behavioral performance by reducing the signal-to-noise ratio of integrated MT neural signals.

Introduction

There is a strong interest in understanding how neural correlations affect visual cortical function. Slow covariations in the activity of neurons (usually measured across trials and referred to as spike-count or noise correlations) can impact the representation

of visual information, and subsequent visually guided behavior (Zohary et al., 1994; Abbott and Dayan, 1999; Sompolinsky et al., 2001; Averbach et al., 2006; Cohen and Newsome, 2008; Cohen and Maunsell, 2009; Ruff and Cohen, 2014; Ni et al., 2018). Importantly, noise correlations between sensory neurons may contribute to the relationship between neuronal firing rate (FR) and behavior (Haefner et al., 2013; Wimmer et al., 2015), and depend on cortical distance and the tuning similarity of the neurons (Ecker et al., 2014; Chelaru and Dragoi, 2016). While the source and impact of noise correlations on coding are still a matter for investigation, reduced noise correlations are linked to enhancements in behavioral performance due to changes in internal state such as attention (Cohen and Maunsell, 2009; Mitchell et al., 2009), expectation (Ruff and Cohen, 2014), adaptation (Gutnisky and Dragoi, 2008), and learning (Gu et al., 2011; Jeanne et al., 2013).

Received April 26, 2017; revised March 23, 2018; accepted March 29, 2018.

Author contributions: A.H. designed research; J.E.T.S. and E.P.C. performed research; A.H., A.G., and E.P.C. analyzed data; A.H., A.G., J.E.T.S., and E.P.C. wrote the paper.

This research was supported by grants from Canadian Institutes of Health Research, the Natural Sciences and Engineering Research Council (E.P.C.), and the Centre for Mathematics in Bioscience and Medicine (A.H.). We thank F. Kingdom for helpful feedback on this work.

The authors declare no competing financial interests.

Correspondence should be addressed to Alireza Hashemi, Department of Physiology, McGill University, 3655 Sir William Osler, Montreal, QC H3G 1Y6, Canada. E-mail: alireza.hashemi@gmail.com.

DOI:10.1523/JNEUROSCI.1182-17.2018

Copyright © 2018 the authors 0270-6474/18/384399-19\$15.00/0

Figure 1. Summary of task, behavioral performance, and multiunit MT responses. **A**, Spike-time correlations are defined as short-timescale correlations in spike activity on a single trial. Illustrated are spikes occurring at about the same time that produce a positive spike-time correlation (left), while spikes occurring asynchronously would result in a negative correlation (right). **B**, Two electrodes recorded simultaneous MU activity in area MT while two monkeys performed a motion detection task. **C**, Schematic of the motion-pulse detection task. On each trial, the monkey fixated on a central cross and two static RDPs were presented. Upon trial start the RDPs began to move with 0% coherence. There was a fixed preamble period of 500 ms where the probability of a motion pulse occurring was zero. After the preamble period, the probability of a motion pulse occurring followed a flat hazard rate. The task of the animal was to release a lever when one or both RDPs contained a 50 ms coherent motion pulse (orange arrows). Thus, the subject did not know when or where (one RDP or both) the motion pulses would occur. The maximum motion-pulse onset time was 10 s. The location, size, motion speed, and direction of each RDP were set to match that preferred by the corresponding MT RF. RFs were nonoverlapping. **D**, After an initial ramp, the time course for the proportion of hits (black) and false alarms (orange) was relatively constant (average of 47 sessions). The proportion of hits was smoothed with a 5 ms Gaussian, while the false alarm rate was computed using 100 ms bins and shifted by the median RT to compensate for the motor delay. **E**, Average population MT MU

In comparison, correlations between neurons on much shorter timescales (measured on a single trial and referred to as synchrony or spike-time correlation) can affect information coding, but their origin and impact on behavior remain a focus of investigation (Singer and Gray, 1995; Steinmetz et al., 2000; Fries et al., 2001a; Kimpo et al., 2003; Thiele and Stoner, 2003; Hirabayashi and Miyashita, 2005; Palanca and DeAngelis, 2005; Mitchell et al., 2009; Huang and Lisberger, 2013; Gomez-Ramirez et al., 2014). For example, attention has been reported to reduce across-trial noise correlations (Cohen and Maunsell, 2009, 2010), but increase spike-time correlations (Steinmetz et al., 2000; Fries et al., 2001a).

Interestingly, recent studies have suggested that the sign of noise correlations (i.e., positive or negative) has implications for readout by downstream brain areas (Downer et al., 2015; Chelaru and Dragoi, 2016). This raises questions regarding the behavioral consequence of the sign of spike-time correlations. Here we investigated how positive and negative spike-time correlations in area MT (middle temporal) are linked to motion detection on a trial-by-trial basis.

Although across-trial measures of noise correlation can be linked to the spike-time correlation (Bair et al., 2001; Huang and Lisberger, 2013), the population of trials used to compute noise correlations likely contains a mix of individual trials with strong positive and negative spike-time correlations. Thus, in principle, the variability of the magnitude of trial-by-trial spike-time correlations could be significant even when the noise correlation is negligible. As illustrated in Figure 1A, spike-time correlations between two sensory pools could vary in sign from trial to trial with no change in spike rate.

We analyzed spike-time correlations in paired multiunit (MU) activity [non-overlapping receptive fields (RFs)] from area MT of monkeys performing a motion detection task. We asked whether these correlations were predictive of behavioral outcome. Compared with single-unit responses, MUs reflect the net activity of a pool of nearby sensory neurons. In our

firing rate for hit trials (blue) and miss trials (red). Responses were aligned either to the start of the 0% coherent motion (left) or the start of the motion pulse (right). The dashed lines are the average response when the motion pulse occurred outside the RF. MU spikes were smoothed with a 10 ms Gaussian. Shading is the SEM.

analysis, we used three simple metrics of spike-time correlation: Pearson's correlation (R), absolute Pearson's correlation ($|R|$), and mutual information (MI).

Just before the motion stimulus occurred, we found that R was a poor predictor of behavioral outcome, while $|R|$ and MI were as good as the firing rate at predicting a correct trial. Positive and negative spike-time correlations were almost equally likely on a given trial, and, once the sign was accounted for, all three measures of neural correlation were similarly predictive of behavioral outcome on a trial-by-trial basis. Our results show that just before the stimulus onset, either strong positive or negative spike-time correlations were associated with more failed detections of the motion stimulus. We propose that both positive and negative spike-time correlations can arise from a common source of noise to the MT pools, which increases the variance of the pool output, and thus, reduces the signal-to-noise ratio (SNR).

Materials and Methods

Behavioral task. The data for this study have been previously described (Smith et al., 2011, 2015; Farah et al., 2014). Two male monkeys (*Macaca mulatta*) were trained to perform a coherent motion pulse detection task outlined in Figure 1C. Stimuli were a pair of nonoverlapping random dot patches (RDPs) with location, size, speed of motion, and direction of motion matched to RF preferences of the recorded neurons. A trial began with the presentation of a fixation point and two static RDPs. Once the monkeys had fixated and pressed the lever, the RDPs remained stationary for an additional 200 ms, upon which dots began moving with 0% coherence (Fig. 1C, trial start).

A 50 ms pulse of coherent motion occurred at a random time from 500 to 10,000 ms in either of the RFs according to an exponential distribution (Fig. 1C, top, flat hazard function). Three possible stimulus conditions were randomly interleaved from trial-to-trial: (1) a motion pulse in RDP 1; (2) a motion pulse in RDP 2; and (3) simultaneous motion pulses in both RDPs. Note that the stimulus was the same (0% coherent) before the motion pulses occurred. After the coherent motion pulse, the RDPs returned to 0% coherent motion. The monkeys had to release the lever while maintaining fixation during a reaction time (RT) window of 200–800 ms after pulse onset (hit trials) to receive a juice reward. The stimulus stopped as soon as the animal released the lever. If the monkey held the lever until the end of the RT window (miss trials), then a final 150 ms of 0% coherent motion was shown before the stimulus stopped and no reward was given. Trials on which the monkey released the lever before the coherent motion pulse (false alarms) were not rewarded. Trials were aborted and not used in our analysis when the monkey did not maintain fixation within 1.5° of the fixation point.

Before training began, animals were implanted with stainless steel posts to stabilize their head position. After training was complete, the animals were implanted with recording chambers (Crist Instruments), and craniotomies were performed to allow a dorsal approach to area MT of visual cortex. Anatomical MRI scans (1.5 T) were performed to verify chamber location and orientation. Surgical procedures were performed in sterile conditions while the animals were anesthetized. Animals received daily care and observation from veterinarians and animal health technicians at the McGill University Animal Care Center. All procedures were approved by the McGill University Animal Care Committee under guidelines set forth by the Canadian Council on Animal Care.

Visual stimulus. Stimuli were presented using a computer monitor placed 57 cm in front of the monkeys (120 Hz refresh, 1600 × 1200 resolution). RDPs consisted of white dots (0.3° wide; density, 10 dots/°²) on a gray background. Dots moved randomly along the preferred-null axis of the receptive fields; during 0% coherent motion, dots had a 0.5 probability of moving in the preferred direction of the neuron independently of other dots. At 100% coherence, all of the dots moved in the preferred direction. Speed was set to that preferred by the neuron. Dots that ran past the edge of the aperture of the RDP were randomly replotted on the opposite side. This design allowed a change in coherence to occur without a change in dot density. Thus, the animals had no cue other than

motion coherence that the motion pulse had occurred. During the motion pulse, the fraction of dots moving coherently was set separately for each RDP to produce threshold performance (~50% correct) in the single motion pulse condition. At the onset of the trial, the dot patterns were at 0% coherence; the duration of this period was determined by a random draw from an exponential function that assumed a flat hazard rate from 500 to 10,000 ms (Fig. 1C, variable motion period). We refer to the period from 0 to 500 ms where there was no chance of a motion pulse as the preamble.

Single-unit recordings. Dual electrophysiological responses from area MT were obtained using standard extracellular recording techniques (Smith et al., 2011). The two electrodes were independently advanced and separated by 1–2 mm (Fig. 1B). After training, the neural responses were collected on 50 sessions. On 47 sessions, we also recorded the raw electrode waveforms, which were used for this current study to extract MU responses. Single neurons were isolated on-line using a dual-window discriminator (Bak), with spike waveforms verified during off-line analysis. After isolating a single neuron, its RF location and size were mapped by hand. The RFs for the two recorded single units were non-overlapping (Smith et al., 2011, their Fig. 1, for the relative location of all RFs). Direction, speed, and size tuning were determined for each isolated unit using the RDP stimulus. The motion detection task was then run using the optimized stimulus parameters as long as isolations could be maintained (number of hit and miss trials collected ranged from 156 to 1389 per session).

Multiunit spikes. To detect MU spikes, we processed our raw electrode recordings (sampled at 25 kHz) using a filter–rectify–filter cascade. Specifically, the raw electrode waveform was first bandpass filtered from 400 to 5000 Hz and then rectified by taking the absolute value. A second low-pass filter at 2000 Hz was applied to the rectified waveforms, and positive-sloped threshold crossings were scored as MU spikes. We used the filter–rectified–filter approach because we wanted to detect spikes that had either maximum positive or negative deflections. Thus, the rectification flipped negative deflections, and the second low-pass filter smoothed the waveforms so that a single threshold crossing was produced for each putative spike. All digital filtering was performed using the zero-phase *filtfilt* Matlab function.

This method of detecting MU spikes is similar to a spike-energy detector. The parameters for our filters were set based on initial visual inspections of the spike waveforms. The threshold was optimized for each recording individually to produce an average 200 spikes/s firing rate in response to the 0% coherent motion just before the motion pulse. This MU spike rate was chosen because it provided reliable spike-time correlation estimates and was ~10× the single-unit spike rate. We compared the single-unit and MU spike responses to our stimulus and observed no qualitative differences [compare Fig. 1E (the MU response), Smith et al., 2011, their Fig. 3 (single-unit response)]. It is possible, however, that our MU spikes contained nonphysiological noise. For example, large simultaneous noise deflections on both electrodes could artificially introduce spike-time correlations. We addressed this by removing trials with the highest electrode variance (see Fig. 9) and found no appreciable effects on our results. In addition, our main findings based on MU spikes were reproduced using our single-unit spike data (see Fig. 6).

Single-trial measures of neural spike-time correlation (R , $|R|$, and MI). Spike-time correlations are defined as the temporal correlation between two spike trains on a single trial. For example, a positive spike-time correlation would arise if the two spike trains tend to produce spikes at nearly the same time (Fig. 1A, left), while a negative spike-time correlation would arise if the spikes occurred asynchronously (Fig. 1A, right). We used three measures to compute single-trial spike-time correlations between the spike activity of our two MT neural pools: R values, $|R|$ values, and MI. The advantages of each method for estimating spike-time correlations is discussed in the Results. Before computing the correlation, spike trains were first smoothed with a Gaussian kernel (Fig. 2A, example trial). Normally a 1.5 ms SD Gaussian was used to smooth spikes, but the effect of Gaussian width on our results is reported in Figures 5 and 6. The width of the Gaussian kernel (as specified by its SD) provides an estimate of the timescale (or frequency range) of the spike-time correlation. Spike-time correlations on a single trial (R , $|R|$, and MI)

between the two smoothed spike trains were calculated using a 300 ms analysis window, with the result aligned to the leading edge of the window (Fig. 2A). The 300 ms window size was set to be long enough to provide enough data for reliable correlation estimates, but still short enough so that we could study the time course of how correlations changed by sliding the analysis window relative to trial start or the onset of the motion pulse.

A detailed discussion of the method we used to calculate MI can be found in the study by Moon et al. (1995). Informally, MI provides a measure of how well we can predict activity in one spike train by looking at the spike density of the other channel. MI provided us with an alternative, well established measure of correlation that produced similar results to $|R|$. In fact, MI and $|R|$ were found to be highly correlated across trials for all multiunit pairs. Although both $|R|$ and MI are positively biased measures of correlation, we were mainly interested in how these measures of correlation differed between hit and miss behavioral outcomes. In addition, we quantitatively linked R (unbiased) and MI (biased) measures of correlation in Figure 6A, which suggested a linear dependence between our two MU recordings.

Linking neural spike-time correlation to behavioral outcome. We used the standard receiver operating characteristic (ROC)-based measure detect probability (DP; Cook and Maunsell, 2002) to quantify the correlation between our measures of spike-time correlation (R , $|R|$, and MI) and behavioral outcome (hit or miss), referred to as DP_R , $DP_{|R|}$, and DP_{MI} , respectively. We also used DP to describe how the average firing rate of our two neural responses was correlated with behavioral outcome within the same analysis window (referred to as DP_{Rate}).

Given that slow fluctuations in firing rate within our 300 ms analysis window could exert an influence on spike-time correlations, we examined these effects in our analyses using a shuffle-subtraction control. Here we shuffled trials (within the same behavioral outcome groups of hit or miss for each session) before computing the spike-time correlation. For each set of shuffle-based spike-time correlations, we computed a single $DP_{shuffled}$ value. This shuffling procedure was repeated 200 times, and the average $DP_{shuffled}$ was then subtracted from the nonshuffled DP value. As shown for $DP_{|R|}$ in Figure 4D, the shuffled $DP_{|R|}$ value tended toward 0.5, suggesting that slow changes in firing rate had little contribution to our DP values. A similar shuffling result was observed for DP_R and DP_{MI} (data not shown). Shuffle subtraction was not applied to our firing rate-based DP_{Rate} .

Model. We constructed a simple model (see Fig. 8H) to explore how positive and negative correlations between our two MT sensory pools could lead to reduced behavioral detection of the motion pulse. The output of each simulated sensory pool was generated using Gaussian-distributed white noise with a 0 mean and a variance of 1. The response of each pool to the motion pulse was modeled by adding a constant to the mean of the white noise. The two pools received a common additive noise input generated

Figure 2. Measures of spike-time correlation and DP for a single session. **A**, Single-trial example for two analysis periods: after the trial began (preamble, left) and before the motion pulse (right). MU spikes were first smoothed by convolving with a 1.5 ms Gaussian kernel. Within a sliding 300 ms analysis window (gray), the smoothed MU responses were used to compute three measures of spike-time correlation: R values (black), $|R|$ values (red), and MI values (blue). The average firing rate for both MU responses was also computed using the same analysis window (green). **B**, Distributions of single-trial spike-time correlations and firing rates for an example session when the analysis window aligned to the start of the motion pulse. Distributions are shown separately for hit (top) and miss (bottom) behavioral outcomes. The corresponding ROC-metric DP is shown above each pair of distributions. DP values that depart from 0.5 suggest a link with behavioral outcome. **C**, Time course of the correlational and firing rate DPs for the same example session. The time corresponds to the leading edge of the 300 ms analysis window for either the preamble period (left) or before the motion pulse (right).

by convolving white noise with a 1.5 ms Gaussian filter (same as the smoothing filter used in our analysis of the neural data). A variable delay between 1 and 25 ms (see Fig. 8H, “D”) was used to jitter the arrival time of the noise reaching one simulated pool to produce both negative and positive correlations. This variable delay was randomly set on each trial.

Without this delay, the model would not produce the range of correlations observed in the data. To vary the magnitude of the correlations in our simulated MU pool output, we adjusted the strength of the common noise input on every trial (a mean of 0.4 and a variance of 0.4 produced correlation values in a range that was similar to our data). Activity from each pool was integrated by convolving with an exponential function with a 50 ms time constant. The model generated 300,000 trials, which we analyzed in the same way as the MU spike data (see Fig. 8I).

Eye position and microsaccade detection. Eye position was sampled at 200 Hz using an infrared tracking system (ASL 6000; Applied Science Laboratories). Because of the noise in the eye tracker, we low-pass filtered (at 20 Hz) the horizontal and vertical eye components during data collection. To detect small eye movements during fixation, the sampled eye signal was linearly interpolated to 1 kHz, and the onset of a putative microsaccade was detected when eye speed crossed an 8°/s threshold. Putative microsaccades were accepted if the resulting amplitude of the saccade was $>0.05^\circ$, and the saccade occurred at least 20 ms after a previous saccade. Note that a portion of the eye data from this experiment has been previously published, and the main sequence of our putative microsaccades produced the expected linear saccadic speed versus amplitude relationship (Herrington et al., 2009, their Fig. 2C). Because our eye signals were low-pass filtered at 20 Hz during data collection, a phase delay was introduced by the filter. Thus, we shifted the saccadic onset times -40 ms to compensate for the phase delay.

Experimental design and statistical analysis. The experimental design is described above and was composed of 47 recordings sessions, which were used for all statistical analyses. Each session had two simultaneous MU recordings.

Statistical significance was computed from *ttest* and *corr* functions in Matlab (MathWorks).

Significance for our DP values were computed by bootstrapping (Efron and Tibshirani, 1986). This was done to account for our single-unit analysis, where weighted averages were computed for DP based on the number of trials (see Fig. 6G,H). Thus, the same bootstrap approach was applied to both single and multiunit datasets. Bootstrapped *p* values were derived from distributions composed of 100,000 bootstrapped samples drawn with replacement. Variability within our data is reported as either SEM or SD, as indicated.

Results

We examined whether trial-by-trial spike-time correlations between two MT MU neural responses predicted the detection by the monkey of a 50 ms coherent motion pulse. Spike-time correlations are defined as the short-timescale covariation of two spike trains on a single trial and can be either positive or negative (Fig. 1A). The single-unit data for this study have previously been used in other studies that addressed the link between firing rate and behavior after the motion pulse occurred (Smith et al., 2011, 2015; Farah et al., 2014). Here we instead focused on the period leading up to the motion pulse. Our goal was to understand how single-trial spike-time correlations between two nonoverlapping sensory pools in area MT (Fig. 1B) evolved over time, and how these correlations were linked to motion detection.

Motion detection task and behavior

Two monkeys were trained to perform the motion detection task illustrated in Figure 1C. At the start of each trial, the animals viewed two static RDPs that began moving with 0% coherent motion at time 0 (trial start). The two RDPs were generated with unique random-number seeds and were thus uncorrelated. We will hereafter refer to the first 500 ms of the 0% coherent motion as the “preamble” (Fig. 1C), during which no motion pulses occurred. The preamble was followed by a variable motion period from 500 to 10,000 ms in which a 50 ms coherent motion pulse could appear in one or both RFs with a flat hazard rate.

The task of the animal was to release a lever within a 200–800 ms RT window after the motion pulse. A lever release within the RT window is referred to as a “hit” outcome, while failure to detect the motion pulse is referred to as a “miss” outcome. A lever release before the motion pulse was scored as a “false alarm.” Because we were interested in the 0% coherent motion period leading up to the motion pulses, we combined trials from the one- or two-pulse conditions in our analyses. Before a motion pulse occurred, all trials were the same (0% coherent motion). Whether one or two motion pulses were presented was randomly interleaved, and the animals did not know in advance where a motion pulse would occur. While the monkeys performed this task, we simultaneously recorded from two recording sites (MU 1 and MU 2) in area MT.

During the preamble, the monkeys knew there was zero probability of a motion pulse. At 500 ms, the probability of a motion pulse was no longer zero and followed a flat hazard rate. This ensured that the animals could not predict when the motion pulse would occur after the initial 500 ms preamble. As such, we observed that after an initial ramp up, the proportion of false alarm and hit behavioral outcomes was relatively constant over time (Fig. 1D). Note that the false alarm rate (Fig. 1D, orange) is shifted leftward by the median RT. These behavioral results suggest that the animals anticipated the end of the preamble and the beginning of the nonzero hazard portion of the trial.

For each experiment, we recorded MT neural activity from nonoverlapping RFs located in the same visual hemifield using two microelectrodes separated by 1–2 mm (Fig. 1B; see Materials and Methods). The location and size of each RDP was matched to those of the RFs, while the direction and speed of the motion pulse were matched to that preferred by the neurons under observation. This increased the probability that our electrodes recorded neural activity that would be correlated with the motion detection task (Bosking and Maunsell, 2011). From previous analysis, the link between our MT single-unit firing rate and behavior on a trial-to-trial basis was relatively strong (Smith et al., 2011).

In 47 sessions, we extracted MU neural population activity from each electrode. The reason we used MU activity was to reduce the detrimental effect that a low number of spikes has on estimating neural correlations, and to ensure that all sessions had approximately the same firing rates. Thus, we set our spike-sorting parameters to produce ~ 200 spikes/s just before the motion pulse occurred (see Materials and Methods) for each of our 94 individual MU recordings. Note that single-unit results are also included below and were similar to our MU results (see Fig. 6).

We first wanted to establish that our MU spikes responded to our stimulus as expected from single-unit MT recordings. Figure 1E (left) shows that the average MU firing rates transiently increased in response to the start of the 0% coherent motion, and then experienced a small decay as the trial progressed, consistent with what is typically observed in MT visual cortical neurons (Müller et al., 2001). MU activity showed a robust response when aligned to the motion pulse (Fig. 1E, right), as expected from our single-unit observations (Smith et al., 2011). Also in agreement with our single-unit observations, the MU responses to the motion pulse were greater for hit trials (solid blue) compared with the miss trials (solid red). Importantly, there was no appreciable MU response when the motion pulse occurred in the other RF (dashed lines), which supports that our RFs were largely nonoverlapping. The difference in firing rate for hit and miss trials after the motion pulse occurred was previously accounted for by

a model with two independent sensory channels (Smith et al., 2011). This model, however, did not include nonsensory effects, such as shifts in attention after the preamble, or the effects of neural correlations. Thus, we next computed three measures of single-trial spike-time correlation between our two MU responses, and linked these measures to behavioral outcome.

The trial-by-trial link between spike-time correlations and behavioral outcome

We hypothesized that single-trial spike-time correlations between our two MU responses would predict behavioral outcome. To examine this hypothesis, we focused on the following two trial periods where the stimulus was 0% coherent: the preamble and just before the motion pulse. We did not include the period after the motion pulse to avoid potential confounds when estimating neuronal correlations during the strong transient increase in spike rate. Thus, our task provided the following two comparison periods: the preamble aligned to the start of the trial, when the animals knew the motion pulse would not occur (probability of a motion pulse = 0); versus the period just before the motion pulse began (probability of a motion pulse > 0).

Figure 2A illustrates how we computed the spike-time correlations between our two MU responses for an example trial. First, MU spikes were smoothed by convolving with a Gaussian kernel (1.5 ms SD), centered on each spike. The width of the Gaussian kernel provides a measure of the timescales (or frequencies) that contribute to the spike-time correlation. As the kernel width increases, the contribution of high frequencies is removed. For example, a 1.5 ms Gaussian kernel has a low-pass –3 dB cutoff of ~88 Hz. We will address the effect of the kernel width below.

Next, spike-time correlations were computed from the smoothed neuronal response within a sliding 300 ms window (Fig. 2A, gray box). A 300 ms window was chosen to provide enough data to reliably estimate correlations while still enabling us to observe their time course. Estimated correlations were aligned to the leading edge of the window (arrows). We stopped the analysis window at 50 ms after the motion pulse because this is when MT typically began responding to the coherent motion.

Using the smoothed MU response, we computed the following three metrics of single-trial spike-time correlation: R , $|R|$, and MI . R is simply the zero-lag correlation between the two smoothed MU responses and is a standard measure of similarity (Cohen and Kohn, 2011; Smith et al., 2011), but has limitations that our other two measures were better equipped to handle. For example, we used the absolute value of the Pearson's correlation because it has been suggested that negative correlations may be as relevant as positive ones (Chelaru and Dragoi, 2016). Mutual information was also used as a measure of single-trial spike-time correlation because it captures nonlinear dependencies and is insensitive to the sign of the correlations. While the single-trial Pearson's correlation is simple to compute, for mutual information we used a common method adapted from the study by Moon et al. (1995) to compute our MI values from density estimates. Note that although $|R|$ and MI values are positively biased measures of correlation, we were interested only in how they differed between hit and miss behavioral outcomes.

Next we linked the spike-time correlation metrics (R , $|R|$, and MI) to behavioral outcome using the ROC-based metric DP (Cook and Maunsell, 2002). DP is similar to choice probability (Britten et al., 1996; Shadlen et al., 1996; Parker and Newsome, 1998; Price and Born, 2010), and is commonly used for expressing the covariation between neural responses and two behavioral outcomes (hit vs miss) on a trial-by-trial basis. A DP of 0.5 indi-

cates that our measure of correlation did not vary with the behavioral performance of the animal. A DP of >0.5 suggests that spike-time correlations were larger on hit trials versus miss trials, whereas a DP of <0.5 indicates the opposite. We estimated the DP at 5 ms time points using the hit and miss distributions of each of our three measures of spike-time correlations and firing rate.

The time course of our three measures of spike-time correlation for the example trial are shown in Figure 2A (R , black; $|R|$, red; MI , blue). The average firing rate across both MU responses within the same 300 ms sliding window is also shown (Fig. 2A, green). Figure 2B illustrates the distributions and corresponding DP values for this example session when the leading edge of the 300 ms window was aligned to the start of the motion pulse. Distributions for hits (top) and misses (bottom) are shown separately for our three measures of spike-time correlation and firing rate.

For the example session shown in Figure 2C, the time course of the DP values was generally >0.5 during the preamble period where the probability of a motion pulse was zero. Thus, spike-time correlations and firing rates were slightly higher for hit trials versus miss trials during this period. By comparison, before the motion pulse the DP values for our three measures of spike-time correlation generally were <0.5, especially $DP_{|R|}$ and DP_{MI} . Thus, trial-by-trial spike-time correlations using the absolute value of the Pearson correlation (Fig. 2C, red) and mutual information (Fig. 2C, blue) were smaller for hit trials versus miss trials just before the motion pulse ($t = 0$), but not for the Pearson's correlation (Fig. 2C, black). As illustrated here, there was usually a significant amount of variability in single-session DPs. In addition, we also computed the DP using the average firing rates for the two MU recordings over the same 300 ms sliding window (Fig. 2C, green). For this example session, the trial-by-trial firing rates were generally higher for hits compared with misses, especially as the trial approached the motion pulse.

Spike-time correlations predicted behavioral outcome just before the motion pulse

The population DP time course of our three measures of spike-time correlation and firing rate are shown in Figure 3A. All three DP values based on correlation (R , $|R|$, and MI) were near 0.5 during the preamble period at the start of the trial. At ~100 ms before motion pulse onset, the population averages of both DP_{MI} and $DP_{|R|}$ (Fig. 3A, blue and red, respectively) began to exhibit the same dynamics with a downward trend that peaked just as the motion pulse began ($t = 0$). The dynamics of DP_R (Fig. 3A, black) by comparison was relatively flat. At the start of the motion pulse, the population means of DP_{MI} and $DP_{|R|}$ were significantly <0.5 ($p = 0.0002$ and $p < 0.0001$, respectively, one-sided bootstrap; $N = 47$), but not that of DP_R ($p = 0.09$). Note that DP values for individual sessions are shown below (Fig. 4). Thus, just as the motion pulse began, the two measures of spike-time correlation that were independent of the sign ($|R|$ and MI) were reliably smaller on hit trials versus miss trials. Toward the end of the 50 ms motion pulse, DP_{MI} and $DP_{|R|}$ values began to rise. The significance of these late dynamics is difficult to interpret, as MT responses begin to be dominated by the coherent motion pulse.

Firing rates, by comparison, had the opposite dynamics for predicting behavioral outcome. The DP_{Rate} (green; Fig. 3A) value was slightly <0.5 during the early preamble, but began an upward trend at the end of the preamble that greatly accelerated during the motion pulse. Interestingly, the magnitude of DP_{Rate} value at the start of the motion pulse ($p = 0.01$, one-sided bootstrap;

Figure 3. Population DP and spike-time correlation time course. **A**, For the three measures of spike-time correlation and spike rate, DPs are aligned to either the start of the preamble period (left) or the start of the motion pulse (right). Firing rate (green) and spike-time correlations computed with $|R|$ (red), MI (blue), but not R (black), were predictive of behavioral outcome just as the motion pulse occurred (time = 0). **B**, Population averages of the three measures of spike-time correlation aligned to the start of the preamble period (left) and the motion pulse (right). Note the change of the ordinate scale between the preamble and motion-pulse alignments. Within each session, trials were grouped according to behavioral outcome: hit (blue) and miss (red). All DPs and averages were computed using the 300 ms sliding window shown in Figure 2A and an MU spike Gaussian smoothing kernel of 1.5 ms. Shading is the SEM.

$N = 47$) was about the same as that of DP_{MI} and $DP_{|R|}$, but in the opposite direction.

This suggests that firing rate and spike-time correlation had similar links to behavioral outcome (equal magnitude, but opposite sign), and is in agreement with other studies that have used across-trial noise correlation measures (Cohen and Newsome, 2008; Cohen and Maunsell, 2011).

Although the number of paired recordings was small ($N = 47$), our experimental design had several beneficial features. First, the stimulus parameters of the motion pulses were matched to that preferred by the neural populations recorded, which has been shown to increase DP values based on firing rate (Bosking and Maunsell, 2011). Second, we collected a large number of trials for each session (range, 156–1389), which helped to reduce the variability in our DP estimates. Third, our long trials and brief 50 ms stimulus placed a premium upon the animals directing their attention to the RDPs at the time that the motion pulses occurred. Once the motion pulse was over, all stimulus information was lost. Finally, during the preamble period at the start of the trial the animals presumably knew that the motion pulses would not occur. Thus, the neural activity during the preamble provides a control period for interpreting DP values. Note that the average time between the end of the preamble and the start of the motion pulse was similar for hit trials and miss trials (781 and 764 ms, respectively).

Our DP values reflected both the differences in the mean of our correlational metrics and their variability. Figure 3B shows the mean values of our three measures of spike-time correlation during the preamble and before the motion pulse. All three mea-

sures show a downward trend during the preamble. This decrease in correlation at the start of a trial has been observed in MT using across-trial measures of correlation (deOliveira et al., 1997; Churchland et al., 2011; Oram, 2011). With the exception of the R values, both $|R|$ and MI values showed an increased separation between hit trials (Fig. 3B, red) and miss trials (Fig. 3B, blue) just before the pulse onset (note the expanded vertical scale for the right column).

The DP values at the start of the motion pulse for individual sessions are shown in Figure 4 (filled triangles are means). The marginal histograms show that for most sessions, DP_{MI} and $DP_{|R|}$ values were <0.5 . We observed a high degree of correlation (Pearson's $\rho = 0.95$) between $DP_{|R|}$ and DP_{MI} values across sessions (Fig. 4A). Thus, $|R|$ and MI were almost equally predictive of trial outcomes. This is in contrast with the lack of correlation observed between $DP_{|R|}$ and DP_R values (Fig. 4B). Furthermore, others have reported a relationship between spike rate and noise correlations (Cohen and Kohn, 2011; de la Rocha et al., 2007). However, we observed no relationship between $DP_{|R|}$ and DP_{Rate} across sessions (Fig. 4C). Finally, we examined whether the slow decreases in firing rates just before the motion pulse (Fig. 1E) contributed to our DP values. We found that $DP_{|R|}$ was unaffected when the effects of slow changes in firing rate were removed from our spike-time correlations using a shuffle-subtraction control (Fig. 4D; see Materials and Methods). Thus, when trials were shuffled, the link between $|R|$ and behavioral outcome was eliminated ($DP_{|R|shuffled} = 0.5$). The same shuffle result was observed for our two other measures of correlation (data not shown).

The DP results so far suggest that just before the motion pulse has reached area MT, two single-trial measures of spike-time

correlation that ignore the sign of the correlation ($|R|$ and MI) were as good at predicting behavioral outcome as firing rate. Thus, both strong positive or negative correlations between the two MU spike trains were more likely to occur when the animals missed the motion pulse. This result raised two questions that we next addressed: what was the timescale of the spike-time correlations that best predicted behavioral outcome? And why was R a poor predictor of behavioral outcome?

The timescale of the spike-time correlations that predict behavioral outcome

We next examined the temporal resolution by which our correlational measures predicted behavioral outcome. For this, we used the analysis window aligned to the start of the motion pulse (Fig. 5, top). First, the MU spike trains were convolved with a Gaussian kernel of variable width, with the SD (σ) of our kernels varying from 0 (no smoothing) to 15 ms (Fig. 5A). The Gaussian kernel provides a notion of the timescale (or frequency range) over which the spike-time correlations were computed. For a kernel width of 0, DP_{MI} and $DP_{|R|}$ values were ~ 0.5 . Gaussian widths between 1.5 and 5 ms generally yielded the strongest DP_{MI} and $DP_{|R|}$ values (note the equivalent low-pass cutoff frequency at the top of Fig. 5A for different kernel widths). At a higher σ value, DP values gradually approached chance levels of 0.5. One possibility for why $DP_{|R|}$ and DP_{MI} values were equal to 0.5 for a zero-width kernel (no smoothing) is that the unfiltered MU spike trains contained high frequencies that masked the low-frequency contributions. As shown, DP_R was generally a poor predictor of behavioral outcome for all Gaussian kernel widths.

Jitter methods have been shown to remove stimulus-locked and slow correlations due to fluctuations in rate (Smith and Somner, 2013). Given that our DPs were strongest for relatively small Gaussian kernel widths, we expected that small amounts of spike jitter would impair the ability of spike-time correlations to predict behavioral outcomes. Thus, DPs were next computed with spikes jittered within a window spanning from 0 to ± 30 ms (using a fixed kernel width of 1.5 ms). Increasing the size of the jitter window beyond a few milliseconds had a marked reduction on DP_{MI} and $DP_{|R|}$ values (Fig. 5B).

To further illustrate the effect of spike jitter, we plot the mean values of our correlational measures as a function of jitter window size. Spike-time correlations (R , $|R|$, and MI) rapidly decreased as a function of the size of the jitter window (Fig. 5C). For $|R|$ and MI, the separation between hit (blue) and miss (red) was highest when there was no jitter. As expected, the link between spike rate and behavioral outcome (DP_{Rate}) was not affected by varying the Gaussian kernel width or jittering spike times (data

Figure 4. Comparison of individual DP values for each session. **A**, $DP_{|R|}$ vs DP_{MI} . **B**, $DP_{|R|}$ vs DP_R . **C**, DP_{Rate} vs $DP_{|R|}$. **D**, $DP_{|R|}$ vs $DP_{|R|}$ minus the shuffle subtraction control (see Materials and Methods). Shuffling trials eliminated the ability of $|R|$ to predict behavioral outcome and produced a $DP_{|R|}^{shuffled}$ value of ~ 0.5 . For all plots, DPs were computed from the 300 ms analysis window aligned to the start of the motion pulse, and MU spikes were smoothed using a 1.5 ms Gaussian. Triangles correspond to mean values, ρ is the Pearson's correlation of the points shown, and dashed lines are the unity slope. The p values are for the marginal distributions computed using a one-sided bootstrap ($N = 47$).

not shown). Overall, these results suggest that correlations on relatively short timescales were linked to behavioral outcome.

The relationship between DP_R , $DP_{|R|}$, and DP_{MI}

Why did DP_R fail to reliably predict behavior? And why were DP_{MI} and $DP_{|R|}$ values similar across sessions? These questions can be addressed by plotting the individual R values versus MI values for all trials where the analysis window was aligned to the start of the motion pulse (Fig. 6A). Although high mutual information does not necessarily imply high correlation values, it is well known that for linearly dependent bivariate Gaussian signals, mutual information is a function of the squared Pearson's correlation coefficient (Pinsker, 1964). As suggested by our data in Figure 6A, our single-trial mutual information values followed the Pearson's spike-time correlation with an exponent of 2.2 (Fig. 6A, red line is fit). This suggests that the trial-by-trial dependencies between our two MU recordings were mostly linear, and that accounting for the negative Pearson's correlation is necessary for predicting a trial outcome.

Figure 6B illustrates the potential confound with using Pearson's spike-time correlation values for computing our ROC met-

Figure 5. Temporal resolution of the spike-time correlation DPs. **A**, DPs as a function of the SD (σ) of the MU spike smoothing Gaussian. $DP_{|R|}$ and DP_{MI} are the strongest for a σ from 1.5 to 5 ms. The corresponding low-pass cutoff frequency for each is shown on top. **B**, Correlational DPs as a function of spike jitter. Each MU spike was randomly moved within a window with or without jitter. Increasing jitter steadily reduced correlational DPs toward 0.5. A jitter = 0 corresponds to the same data as in Figure 4. A smoothing Gaussian of 1.5 ms was used during the jitter analysis. **C**, Population means for the three measures of spike-time correlations as a function of jitter for hit (blue) and miss (red) trials. All plots used the same analysis window aligned to the start of the motion pulse. Shaded area is the SEM.

ric DP_R . If we assume that missed trials were associated with both stronger positive and negative Pearson's correlations, then this would cause DP values to be pulled toward 0.5 (all trials, left). If this were the case, then examining trials with only positive correlations (middle) or negative correlations (right) would yield DP_R values that depart from 0.5 in opposite directions. Thus, for trials with a positive Pearson's correlation this hypothesis would predict DP_R values <0.5 ; whereas for trials with negative spike-time correlations, DP_R values would be >0.5 .

To test this model, we recomputed DP values grouped on the sign of the Pearson's spike-time correlation (Fig. 6C, analysis window aligned to the start of the motion pulse). There was about an equal number of trials with positive and negative correlations on each session (average $\pm$ SD proportion of trials with $R > 0$, $51.7 \pm 6.7\%$, $N = 47$; Fig. 7B). Although the mean population DP_{MI} (Fig. 6C, blue) and DP_{Rate} (Fig. 6C, green) values remained relatively the same for trials with only positive or negative correlations, DP_R (Fig. 6C, black) values became stronger and deviated significantly from 0.5 in opposite directions depending on the sign of the correlation ($DP_R < 0.5$ for $R > 0$ and $DP_R > 0.5$ for $R < 0$). The p value in Figure 6C reports the pairwise difference in DP_R between

positive and negative correlations (one-sided bootstrap, $N = 47$). This result confirms that DP_R was a poor ROC-based predictor of trial outcome due to the combined effect of both stronger positive and negative correlations associated with missed trials.

Individual session DP_R values are shown in Figure 6D for each group ($R > 0$, open circles; $R < 0$, black circles), and they followed DP_{MI} as suggested by the approximately linear dependence between R and MI highlighted in Figure 6A. The p values in Figure 6D correspond to DP_R values either >0.5 or <0.5 (one-sided bootstrap, $N = 47$). With the exception of the magnitude relative to 0.5, the effect of the smoothing kernel width (σ) on DP_R was similar for trials with negative and positive R values (Fig. 6E). Additionally, we found that behavioral outcomes were, on average, the same regardless of the presence of either negative or positive neural correlations (Fig. 6F). Thus, the link between neural correlations and behavior appears to be the same regardless of the sign of the spike-time correlation. These observations account for the differences and similarities among DP_R , $DP_{|R|}$, and DP_{MI} values.

Finally, we sought to evaluate whether the results obtained using MU responses were in agreement with our single-unit re-

Figure 6. The relationship between DP_R , $DP_{|R|}$, and DP_{MI} is dependent upon the sign of the spike-time correlation. All panels show results for the analysis window aligned to the start of the motion pulse (top). **A**, Single-trial MI and spike-time R values computed from the smoothed MU responses ($\sigma = 1.5$ ms) for the entire dataset. Note that only every other data point was plotted for clarity. Single-trial values of MI followed R values with an exponent of 2.2 (red line is fit to all data points). **B**, Model for how an R value might affect the ROC-based DP metric. If single-trial values of R are both positively and negatively stronger on missed trials, then for all trials $DP_R = 0.5$ (left). In comparison, trials with $R > 0$ would produce $DP_R < 0.5$ (middle), while trials with $R < 0$ would produce $DP_R > 0.5$ (right). **C**, Summary of population DP values using all trials (left bar plots; same data from Fig. 4), trials with $R > 0$ (middle bar plots), and trials with $R < 0$ (right bar plots). When trials were grouped by the sign of R , DP_R (black bars) values became larger (i.e., departed from 0.5) and reversed sign, as predicted by the model. There was an approximately equal number of trials with positive and negative R values. **D**, Individual DP_R and DP_{MI} values from all sessions for trials with $R > 0$ (open) and $R < 0$ (black). Triangles are the means. The p values are one-sided tests for mean $DP_R > 0.5$ ($R < 0$) and mean $DP_R < 0.5$ ($R > 0$), bootstrap ($N = 47$). **E**, DP_R as a function of the SD (σ) of the MU spike smoothing Gaussian for trials with $R > 0$ (dashed line) and $R < 0$ (solid line). **F**, Behavioral performance (proportion hits) for trials with negative spike-time R values vs positive spike-time correlations. The dashed line is unity slope. **G**, Summary of DPs using single-unit spikes, grouped by the sign of the spike-time correlations. Note that the single-unit spikes were convolved with a 4 ms Gaussian, and a weighted average based on the number of trials was used to compute the mean DPs (see Results). Only trials with at least two single-unit spikes per electrode within the analysis window were included. **H**, Single-unit DP_R as a function of the SD (σ) of the spike-smoothing Gaussian for trials with $R > 0$ (dashed line) and $R < 0$ (solid line). Error bars and shading are SEM, and p values are from paired one-sided bootstraps that DP_R ($R < 0$) is greater than DP_R ($R > 0$; $N = 47$).

cordings. It has been suggested that spike sorting can increase or decrease across-trial measures of correlation (Cohen and Kohn, 2011). Estimating single-unit spike-time correlations was difficult because some MT neurons had relatively low firing rates. Thus, an estimate of the single-trial spike-time correlation required a minimum of at least two spikes from each single unit to fall within the 300 ms analysis window. Trials with only a single spike or no spike in the analysis window were excluded. We also used a slightly wider Gaussian kernel width of 4 ms to smooth the spikes to reduce the effects of the low firing rates. To compensate for the reduced number of trials for some single units, the population DP values shown in Figure 6G are a weighted average, where the weighting was based on the total number of trials used to compute the DP. Although the single-unit DP values were

weaker than the MU data for our three correlational measures and firing rate, they reproduced the same trends when grouped by positive or negative correlations. Notably, single-unit DP_R (Fig. 6G, black bars) showed a significant pairwise flip from <0.5 to >0.5 when separated by the sign of the correlations (one-sided bootstrap, $N = 47$). Across different smoothing kernel widths, the single-unit DP_R had a similar shape as that for the MU data (compare Fig. 6E,H).

Detection performance versus single-trial spike-time correlations (R)

We have shown that when the sign was accounted for, the single-trial Pearson's correlation of the two Gaussian smoothed MU responses before the motion pulse predicted behavior in our de-

Figure 7. Spike-time R values and FRs modulate behavioral performance. **A**, Example single-trial R vs FR for hit (blue) and miss (red) trials from an example session (ρ is the Pearson's correlation of all data points). **B**, Bottom, Distribution of the Pearson's correlation between R and FR (ρ). Top, Distribution of the proportion of trials with $R > 0$. Histograms include all sessions ($N = 47$), and triangles are the means. **C**, Raw and normalized behavioral performance binned on R and FR for an example session. Behavioral performance was normalized by (Figure legend continues.)

tection task. What is not revealed by the DP metrics used above, however, is the magnitude of this effect or a potential mechanism. For example, why would both positive and negative correlations before the stimulus lead to reduced detection performance? And how might both positive and negative spike-time correlations arise?

One important aspect of our data was the large variability in the single-trial spike-time R value within each recording session. The scatter plot in Figure 7A illustrates this variability for a single session for hit trials and miss trials (note the analysis window was aligned to the onset of the motion pulse). It has been reported that measures of correlation increase with firing rate (de la Rocha et al., 2007; Cohen and Kohn, 2011). As illustrated by the example session in Figure 7A, we did not observe this relationship between R and FR (Pearson's correlation $\rho = 0.002$ using all trials). Across sessions, the average $\pm$ SD Pearson's correlation between R and FR was 0.02 ± 0.09 ($N = 47$), as shown by the histogram in Figure 7B (bottom). As mentioned above, we also typically observed an almost equal number of positive and negative single-trial values of R within a session (Fig. 7B, top).

To illustrate the magnitude of how detection performance changed as a function of single-trial spike-time R values and FRs, we binned trials on either R or FR, and computed the proportion of hits (i.e., correct outcomes) within each bin (Fig. 7C, example session). Note that bin edges were chosen to produce an approximately equal number of trials per bin and that the middle correlational bin was centered on $r = 0$. In the example session in Figure 7C, the highest proportion of hits was associated with values of R near zero and high FRs. For each session, we normalized how behavior was modulated by R and FR by dividing by the mean. For the example session shown, hits were modulated by approximately $\pm 6\%$ about the mean (12% total modulation) as a function of R . FR by comparison was associated with a $\pm 10\%$ modulation of hits about the mean.

The average modulation of the proportion of hits is shown in Figure 7D as a function of R and FR. Both the strongest positive (Fig. 7D, green) and negative (Fig. 7D, blue) correlations significantly reduced detection performance by $\sim 9\%$ when compared with trials with little or no correlation (Fig. 7D, yellow). The p values in Figure 7D are paired comparisons (one-sided bootstrap, $N = 47$). As shown on the right, there was a similar modulation in detection performance between the smallest and largest FRs.

To further illustrate how spike-time R and FR values jointly contributed to behavioral outcome, we binned trials according to both their firing rate and correlation values (25 bins), and examined the proportion of hits within each bin. Because of the large number of bins, the joint contribution of R and FR on detection performance was relatively noisy for a single session (Fig. 7E). However, the population average across sessions revealed that high FRs and low correlations were associated with more hits (Fig. 7F). In comparison, the fewest hits occurred in bins associated with low firing rates and high positive and negative correlations.

Both positive and negative spike-time correlations are associated with increased variance. To examine potential mechanisms for the modulation of hits as a function of R , we first confirmed that the average firing rate for each correlation bin had the same dynamics (Fig. 7G, colors correspond to the same points in D). As shown, there was no effect on the average MU time course, before or after the motion pulse, when trials were binned by the spike-time R value (computed in the 300 ms before the motion pulse). Note that only trials where a motion pulse occurred in the RF of the MU were included. Figure 7H summarizes this result by plotting the average normalized firing rate as a function of R using the two analysis windows shown in Figure 7G (before and after pulse, black bars). Thus, the mean MU response does not seem to explain the modulation in behavioral performance associated with spike-time correlations.

The variance of sensory information is commonly thought of as a major factor during decision-making (Churchland et al., 2011; Zylberberg et al., 2016). If downstream brain areas are integrating MT population activity to detect the weak motion pulses, then the amount of variance in the integrated response before the pulse becomes important. For example, we simulated the integration of our MU activity by convolving spikes with an exponential kernel ($\tau = 50$ ms; Fig. 8A). The integrated MU response after the motion pulse would likely be more salient when preceded by activity with less variance (Fig. 8A, low variance, left). Likewise, a higher variance of the integrated response before the motion pulse could make the response after the pulse harder to detect (Fig. 8A, high variance, right).

Thus, we examined the variance of our integrated MU responses as a function of the spike-time R value before the motion pulse (note that the R value was binned the same as in Fig. 7).

As illustrated by the example session, we commonly observed that the variance of the integrated MU responses before the motion pulse (Fig. 8A, σ^2) was less for values of R near zero (Fig. 8B). Note that this is the same example session shown in Figure 7 with the same binning on R . The population-normalized variance of the integrated response as a function of R (Fig. 8E, filled circles) showed a similar shape. In comparison, the variance of the integrated MU response after the motion pulse was relatively flat as a function of R (Fig. 8E, open squares).

As shown by the example session, the mean integrated MU response after the motion pulse (Fig. 8A, μ), was only weakly modulated by spike-time correlations (Fig. 8C), especially for the population average (Fig. 8F, open squares). The mean population response before the motion pulse was also a relatively flat function of R (Fig. 8F, filled circles).

The variance of the integrated response before the motion pulse suggests a way to link spike-time R values to the modulation in behavioral hits shown in Figure 7D. We quantified the ability of the integrated MU response to signal the occurrence of motion using an SNR measure that was a function of the mean integrated response after the pulse (Fig. 8A, μ) and the variance of the inte-

(Figure legend continued.) the mean (dashed line). **D**, Average of the normalized behavioral performance (proportion hits) binned separately for R and FR. For each session, we first normalized the behavioral performance by dividing by the mean. R and FR bin centers are the average of the centers across all sessions. Bin centers were chosen separately for each session so that approximately the same number of trials contributed to each bin. The p values are for the paired comparisons shown, one-sided bootstrap ($N = 47$). **E**, Behavioral performance jointly binned on both R and FR values for the example session in **A**. Data were binned using a 5×5 grid (25 bins total), and bin edges were selected so that each bin had approximately the same number of trials. **F**, Population average of the normalized behavioral performance jointly binned on R and FR values. Bin centers reflect the average across all sessions, and each session was normalized by the mean proportion of hits. Hits were more likely on trials with both higher firing rates and spike-time correlations near 0. **G**, Average MU firing rates for trials binned by R . Colors correspond to the same bins in **D**. Note that all five traces significantly overlap and that only trials where the motion pulse occurred in the RF of the MU were included. Representative error bars are shown for clarity. **H**, Average normalized MU firing rate as a function of R . Normalization was accomplished by dividing single-trial firing rates by the mean firing rate for that session. Firing rates were computed either before the pulse (filled circles and solid line) or right after the pulse (open square and dashed line). Black bars in **G** indicate the time window used to compute firing rates. Error bars are the SEM and, for some, plots are smaller than the symbols.

Figure 8. Positive and negative single-trial spike-time R values are associated with increased variance in the neural response. **A**, Two example trials where the MU spikes were integrated with an exponential filter (time constant, 50 ms). Low variance of the integrated response (σ^2) before the motion pulse could make the response to the motion pulse more salient (left). In contrast, the high variance of the integrated response before the motion pulse could have the opposite effect (right). The effect of variance before the motion pulse can be captured by defining an SNR as the ratio of the mean integrated response after the motion pulse (μ) to the variance of the integrated response before the motion response (σ^2). **B**, The average modulation in the variance of the integrated MU response before the motion pulse as a function of R for the example session in Figure 7. Shown is the variance of one electrode. The y-axis shows both the raw variance (left) and the variance normalized to the mean (right). The analysis window was aligned to the start of the motion pulse, and the MU smoothing kernel was $\sigma = 1.5$ ms. For all plots in this figure, R was binned in the same manner as shown in Figure 7. **C**, Example of the average modulation in the mean integrated response after the motion pulse as a function of R . Note that only trials where the motion pulse occurred in the RF of the MU were included. **D**, Example of the average modulation of SNR as a function of R . SNR was computed for each trial as shown in **A**, and then averaged across trials for each session. **E**, Population average modulation of the normalized variance of the integrated MU response as a function of R . Shown is the variance computed before and after the motion pulse. **F**, Population average modulation in the normalized mean integrated response as a function of R . Shown is the mean response computed before and after the motion pulse. **G**, Population average modulation of the normalized SNR as a function of R computed for a range of integrator time constants (5–100 ms). Only error bars for $\tau = 50$ are shown for clarity. For **E–G**, each of the 47 sessions contributed two MU recordings. R is the Pearson's correlation of the MU response computed in the 300 ms window before the motion pulse (same data as in Fig. 5). **H**, Functional model of the two MT sensory pools. Simulated MU activity was generated with a Gaussian distributed random variable where the mean increased in response to the motion pulse (for details, (Figure legend continues).)

grated response before the pulse (Fig. 8A, σ^2). Figure 8D shows that the average SNR for the example session is modulated by R , such that the highest SNR corresponds to the smallest values of R . Across our population, the average normalized SNR of our integrated MU responses (Fig. 8G, black) was modulated by R in a similar manner as the proportion of hits (compare Fig. 7D). Note that only trials where the motion pulse occurred in the RF of the MU were used to compute SNR.

For every session, the SNR was computed for each trial individually and then averaged across trials for each R bin (Fig. 8D). One potential problem with this approach was that trials with very low variances could produce exceedingly high SNR outliers that bias the average session SNR. We also computed the median SNR for each session and observed the same SNR shape as shown in Figure 8G (data not shown). Thus, our SNR values were not biased by extreme outliers.

As we do not know how downstream areas integrate MT responses, we computed SNRs for a range of integration time constants (Fig. 8G). In all cases, the highest population SNR corresponded to spike-time R values near 0; however, the magnitude of the SNR modulation decreased for smaller integration time constants.

It has been suggested that neural correlations may stem from global fluctuations of neural activity (Ecker et al., 2014). Thus, a common input that modulated our two MT neural pools could lead to both increased spike-time correlations and variance. In this scheme, strong spike-time correlations (either positive or negative) would be associated with an increase in the variance of the integrated MT activity, making it more difficult for downstream areas to detect the brief motion-pulse transient.

To examine this hypothesis, we constructed a simple computer model based on random Gaussian distributed signals to simulate our two MT MU responses (Fig. 8H; see Materials and Methods). A common additive noise input created spike-time correlations between the two simulated MU responses. This common noise alone, however, was unable to produce the same range of negative spike-time correlations observed in our data. To mimic both negative and positive correlations observed in our data required adding a small amount of temporal jitter (D in Fig. 8H, 1–25 ms) to the common input arriving at one simulated MT pool. As shown in Figure 8I, across 300,000 simulated trials, this produced variance and SNR plots of the integrated response of the model as a function of R that were similar in shape to that of our data.

Our model suggests that variability in the relative phase of the common noise arriving at the two MT pools is sufficient to produce negative correlations and add variance to the output of each MT sensory pool. For example, variability in the relative phase of the two noise inputs could arise if the animals were dynamically moving their attentional focus between the two stimuli. However, we can only speculate as to whether the variable delay in the model is a physiologically realistic mechanism.

Other contributions to the link between spike-time correlations and behavior

Last, we examined the following three potential contributions to the trial-by-trial spike-time correlations in our MU responses and their corresponding link to behavioral outcome: (1) fluctuations in our random dot stimulus; (2) small eye movements; and (3) the variance of our electrode signals. Just before the motion pulse occurred, the two RDPs contained independent 0% coherent motion. Although the 0% coherent stimulus was designed to have no net motion, small stochastic fluctuations were present in the proportion of dots moving in either the preferred or null directions. We first asked whether these small fluctuations contributed to the link between MU spike-time correlations and behavioral outcome.

In 36 of 47 sessions, we had saved the dot pattern for every stimulus frame. From these sessions, we were able to estimate the frame-by-frame motion strength by computing the net number of dots that moved in either the preferred or the null direction (similar to Cook and Maunsell, 2004). Figure 9A shows a single-trial example of the frame-by-frame net motion aligned to the pulse onset ($t = 0$), where values >0 indicate a net preferred-direction motion and values <0 represent a net null-direction motion. For each RDP, we normalized the motion during the 0% coherence period to have an SD of 1 across all trials (Fig. 9A, note scale bar).

To assess our characterization of the RDP motion, we computed the MU spike-triggered average (STA) of the motion during the 0% coherence. As shown by the example STAs from a single session (Fig. 9B), MU spikes tended to follow motion in the preferred direction with a latency of ~ 50 ms (Fig. 9B, black). In addition, we found that STAs were flat (Fig. 9B, orange) when computed using the motion in the other RDP located outside the RF. Our population average STA (Fig. 9C) suggests that our MU spikes were correlated with the frame-by-frame motion description and that there was little contribution from the motion in the RDP located outside the RF.

We next examined whether correlations in the frame-by-frame motion stimulus (R motion) were associated with MU spike-time correlations (R). Note that the R value is from the same data presented above and that the analysis window was aligned to the start of the motion pulse. The window for the motion stimulus was shifted 50 ms to account for the MT latency observed in our STAs (Fig. 9A, gray). As shown in Figure 9D, correlations in the motion did not seem to be related to spike-time correlations. Furthermore, to determine whether the correlations in the frame-by-frame motion predicted behavioral outcome, we computed the $DP_{|R|}$ value using the motion stimulus and compared it with $DP_{|R|}$ value computed from the MU responses (Fig. 9E). Although our subpopulation of 36 sessions still showed most MU $DP_{|R|}$ values <0.5 (Fig. 9E, triangles are the mean), the mean $DP_{|R|}$ value computed from the frame-by-frame motion stimulus was not significantly different from 0.5 ($p = 0.33$, by bootstrap; $N = 36$). Note that there was no significant correlation between the MU and motion DP values (Pearson's $\rho = 0.21$, $p = 0.17$). Although the STAs show that our MU responses were linked to the stochastic fluctuation in the 0% coherent motion stimulus, these stimulus-driven fluctuations did not appear to account for the link between MU spike-time correlations and behavioral outcome.

Care must be taken when interpreting this result, however, as we cannot rule out all aspects of the RDP motion as a potential contribution. For example, local dot interactions within the RF that would not be captured by our frame-by-frame global motion

(Figure legend continued.) see Materials and Methods). The simulated MU was then integrated with a 50 ms exponential. An additive common noise input (also modeled as a Gaussian random variable) was applied to the output of each sensory pool. To produce both positive and negative spike-time correlations of the simulated MU response, a variable delay changed the relative phase of the common noise arriving at one pool on every trial between 1 and 25 ms. *I*, The variance of the model and the SNR of the integrated response as a function of R of the simulated MU activity (average of 300,000 trials). Note that the R value of the simulated MU of the model was computed in the same manner as the recorded MU activity.

Figure 9. The contributions of the stimulus, small eye movements, and electrode variance to the link between single-trial spike-time correlations values and behavioral outcomes. **A**, In 36 sessions, we were able to extract the dot pattern of the RDP stimulus. Example random fluctuations in the motion for each RDP. Motion was described as the net number of dots moving in the preferred or null direction (see Materials and Methods) and normalized by the session SD. **B**, Example MU STAs of the motion from a single session. Orange is the STA using the motion in the other RDP. **C**, Population average MU STAs. **D**, The correlation of the motion in the two RDPs (R for motion) vs the MU spike-time R value. **E**, $DP_{|R|}$ was also computed using the RDP motion instead of the MU spikes. Across sessions, the average motion-based $DP_{|R|}$ value was near 0.5 and not strongly correlated with $DP_{|R|}$ values computed from the MU spikes (ρ). **F**, Example trial with three putative microsaccades (*) within the analysis window (gray) just before the motion pulse. **G**, Population microsaccade-triggered average of MU activity computed during the 0% coherent motion (orange is the shuffled control of microsaccade times). Note that the microsaccade onset time was shifted -40 ms to account for the 20 Hz low-pass filtering of the recorded eye signals (see Materials and Methods). **H**, The absolute value of the MU spike-time correlations (i.e., $|R|$) vs the number of microsaccades for hit trials and miss trials. A total of 25% of trials had putative microsaccades within the analysis window just before the motion pulse. **I**, Normalized microsaccade rate vs MU spike-time R values. **J**, $DP_{|R|}$ computed with the MU spikes with and without microsaccades. (Figure legend continues.)

description could potentially contribute to spike-time correlations. Another possibility is that chance-correlated movement between the dots on the nearest edges of the two RDPs could simultaneously contribute to both MT pools. Although not used here, one way to better measure the stimulus contributions to spike-time correlation would be to include repeated stimulus dot patterns and examine whether correlations vary under identical 0% motion (Wimmer et al., 2015).

Next, we examined the effects of small eye movements on our MU responses. During a trial, small eye movements such as microsaccades can modulate neural activity in area MT (Herrington et al., 2009) and areas that drive MT such as V1 (Snodderly et al., 2001; McFarland et al., 2016), which could introduce spike-time correlations between our two MT sensory pools. We extracted the occurrence of putative microsaccades as shown by the example trial in Figure 9F (asterisks indicate that three saccades in the 300 ms analysis window aligned the motion pulse; see Materials and Methods). Note that due to the limitations of our eye-tracking system, eye signals were low-pass filtered at 20 Hz, which limited our ability to detect high-frequency saccades (see Materials and Methods). We assessed the relationship between our putative microsaccades and neural activity by computing a saccade-triggered average of the MU spikes (Fig. 9G). Consistent with previous results (Martinez-Conde et al., 2000), we observed a significant modulation in firing rates around the time of these small eye movements (Fig. 9G, orange is the control MU firing rate aligned to randomly shuffled microsaccades).

Given that microsaccades modulated the response of our MU recordings, we also assessed the relationship between the number of microsaccades and the absolute value of the MU spike-time correlation ($|R|$) in the 300 ms analysis window aligned to the motion pulse (Fig. 9H). We detected putative microsaccades within this window on 25% of trials, and these trials tended to have higher values of MU $|R|$ for both hit (Fig. 9H, blue) and miss (Fig. 9H, red) outcomes. Interestingly, the rate of microsaccades was associated with increased positive spike-time R values, but not negative R values (Fig. 9I).

Figure 9, H and I, raised the possibility that microsaccades contributed to the spike-time correlations in our MU recordings. However, when the 25% of trials containing putative microsaccades were removed from the analysis, $DP_{|R|}$ values remained relatively unaffected (Fig. 9J).

Furthermore, eye-velocity variance had the same effect on both MU spike-time correlations and $DP_{|R|}$ values as our putative microsaccades (data not shown). It is possible that our microsaccade detection may have missed fast eye movements. Thus, we cannot fully rule out the potential contribution of small eye movements to the link between spike-time correlations and behavioral outcome.

Finally, we sought to address the possibility that large recording anomalies that occurred simultaneously in both electrodes contributed to the relationship between spike-time correlations and trial outcome. We hypothesized that if nonphysiological sources of noise produced large fluctuations in the variance of both electrode waveforms (Fig. 9K, example electrode recordings), then removing these high-variance trials would result in a weaker $DP_{|R|}$. Using the 300 ms analysis window aligned to the motion pulse (Fig. 9K, gray) and binning the average electrode variance with an equal number of trials per bin, we observed a weak but positive relationship between the average electrode variance and MU $|R|$ for both hit (Fig. 9L, blue) and miss (Fig. 9L, red) trials. The difference between hit and miss MU $|R|$ values, however, became smaller as electrode variance increased. Unsurprisingly, we also observed that electrode variance increased for both positive and negative spike-time correlations (Fig. 9M). However, $DP_{|R|}$ values were relatively unchanged when the 33% highest-variance trials were removed (Fig. 9N). This suggests that trials with large electrode anomalies were not primary contributors to the link between MU spike-time correlations and behavioral outcome.

Discussion

Combining paired recordings with a detection task, we examined the link between positive and negative MU spike-time correlations and behavioral outcomes on a trial-by-trial basis. We reasoned that if neural correlations have an effect on behavioral outcomes, then the correlations present on a single trial ought to be more informative than across-trial measures (e.g., noise correlations). Specifically, we focused on spike-time correlations between two MT sensory pools with nonoverlapping RFs (electrodes separated by 1–2 mm) just before the motion pulse occurred. Within a session, single-trial positive or negative spike-time correlations were about equally likely, and once their sign was accounted for, our three measures of spike-time correlations (R , $|R|$, and MI) produced similar links to behavior. Correlations on short timescales (<5 ms) were as good as firing rates at predicting behavioral outcome. Importantly, the presence of either positive or negative spike-time correlations was associated with a failure to detect the motion pulse.

Just before the motion pulse, the sign and magnitude of spike-time correlations varied greatly from trial to trial and were not appreciably dependent on firing rates. Because of the similar number of trials with positive or negative correlations, the correlational sign was not likely to be due to a systematic relationship between the two MU recordings, which would tend to produce the same correlational sign on all trials. However, we do not know the source of the variability in our single-trial spike-time correlations. Nevertheless, we found that strong positive and negative spike-time correlations were associated with an increase in the variance of the integrated MU response. We propose that this increase in neural variance leads to a decrease in SNR and a subsequent reduction in behavioral performance.

Both intracellular and extracellular recordings suggest that coordinated presynaptic spikes arriving within a short window are more likely to produce a postsynaptic response (Gasparini et al., 2004; Losonczy and Magee, 2006; Zandvakili and Kohn, 2015). Spike-time correlations on short timescales are also implicated in improvements in visual processing. For example, spike synchrony is thought to enhance the response of downstream neurons (Singer and Gray, 1995; Steinmetz et al., 2000; but see Palanca and DeAngelis, 2005; Womelsdorf and Fries, 2007; Fries, 2009). Increases in spike-spike, LFP and LFP-spike gamma band

(Figure legend continued.) Removing the 25% of trials containing microsaccades had little effect on $DP_{|R|}$. **K**, Electrode waveforms for an example trial. Waveforms were normalized by their session SD. **L**, MU $|R|$ as a function of the binned average electrode variance for hit trials and miss trials. Bin edges were set on each session so that an equal number of trials contributed to each bin. **M**, Normalized electrode variance vs MU spike-time correlations. **N**, $DP_{|R|}$ computed with the MU spikes with and without the highest electrode-variance trials. Removing the 33% of trials with the highest variance had little effect on $DP_{|R|}$ values. Note that in all plots, data are from an analysis window aligned to the start of the motion pulse, except for the motion description (**A**), which was shifted by -50 ms to accommodate for the latency of MT. Also note that the MU $DP_{|R|}$ values for all trials are the same data as in Figure 5. Error bars and shaded areas are the SEM. The p values in **E**, **J**, and **N** are for the mean DP values <0.5 ; one-sided bootstrap and small triangles are the mean DP values.

coherence (between 30 and 80 Hz) are known to be modulated by changes in attention and working memory (deOliveira et al., 1997; Fries et al., 2001a, 2008; Pesaran et al., 2002; Bosman et al., 2012). In contrast, a few studies report the detrimental effect of spike-time correlations for low frequencies (Mitchell et al., 2009; Herrero et al., 2013). The role of the sign of the spike-time correlations on behavior, however, has received much less attention.

The contribution of spike-time correlations to performance could be due to the nature of our task. Our two MU pools were relatively distant with nonoverlapping RFs, and a previous analysis of these data suggested that the animals treated each MT pool independently (Smith et al., 2011). Given that we used a very brief 50 ms stimulus, neural activity in the interval directly preceding motion pulse onset was likely to have had a significant impact on behavioral outcomes. Once the motion pulse ended, there was no available sensory information to capture the attention of the animal, as may be the case with longer stimulus exposures (Reynolds et al., 2000; Lee and Maunsell, 2010). Thus, our task design placed a premium on the attentional state of the animal directly before the 50 ms motion pulse. The eventual location of the actual motion pulse did not affect our results. For example, the $DP_{|R|}$ value was the same when computed using trials where a motion pulse occurred in either one or both RDPs (data not shown).

The proposed mechanism for our results is a common noise source to both MT pools, which may represent attentional state. However, spike-time correlations can also be induced by other mechanisms. For example, Huang and Lisberger (2013) have shown that short-timescale correlations between nearby MT neurons (~ 0.3 mm separation) can be modeled by circuit interactions. We did not find similar interactions in our MU data, presumably because our MU pools were further apart. Many studies have suggested that sensory information is integrated back in time until a decision is made (Ditterich et al., 2003; Smith and Ratcliff, 2004; Huk and Shadlen, 2005; Gold and Shadlen, 2007; Smith et al., 2011; but see Katz et al., 2016). In this model, the neural variance becomes an important factor driving behavioral responses (Churchland et al., 2011; Zylberberg et al., 2016). The variance of our integrated MU responses before the motion pulse was associated with spike-time R values. Similarly, others have proposed that trial-by-trial fluctuations in attention reduce the variance of a shared modulator and lower correlation values, which in turn reduces the summed variance across the population (Rabinowitz et al., 2015). Thus, we captured the relationship between variance and R with a functional model where spike-time correlations were driven by an additive common noise source with a variable delay to each MT pool. We can only speculate how a variable delay might arise in the cortical circuitry. Although traveling waves could, in theory, produce temporal delays in the neural response (Sato et al., 2012), we are hesitant to prescribe a specific mechanism for this functional model.

Negative correlations have been reported in visual cortex (Gutnisky and Dragoi, 2008; Ecker et al., 2010; Jeanne et al., 2013), but their functional role is a matter of investigation and is thought to depend on their relationship to the tuning properties of individual neurons (Averbeck et al., 2006; Kohn et al., 2016; Latham and Roudi, 2011). For example, recent studies have reported an impact of negative correlations on population coding (Jeanne et al., 2013; Chelaru and Dragoi, 2016). Previous studies have also shown that the sign and magnitude of noise and spike-time correlations is tied to their tuning similarity and cortical distance (Fries et al., 2001b; Ecker et al., 2014; Ruff and Cohen, 2014). However, we found no relationship between the RF properties of our MU pools and the strength of the spike-time corre-

lations in our data. By the time the motion pulse occurred, our R values were very weak (Fig. 3B). This, along with the large 1–2 mm distance separating our electrodes, may have contributed to the lack of an observable relationship between correlations and RF properties.

Given that fluctuations in the 0% coherence stimulus did not predict behavioral outcomes (Fig. 9D), it is possible that neural correlations just before the motion pulse reflected the attentional state of the monkey (Kohn et al., 2009). Recent studies have shown that shared activity in populations of sensory neurons is modulated by trial-to-trial fluctuations in attentional state (Denfield et al., 2017). Although cues are explicitly provided in most attention studies, shifts in attention still occur (Pang et al., 1992). For example, a subject's prior knowledge about the temporal structure of a task may significantly modulate neural activity and behavioral outcomes (Ghose and Maunsell, 2002; Coull, 2004; Wright and Fitzgerald, 2004; Doherty et al., 2005; Shuler and Bear, 2006). In addition, the onset of a sensory stimulus has been previously shown to produce a decline in neuronal variability and covariability in MT (Churchland et al., 2010; Oram, 2011), and would account for the initial decline in correlations observed during our preamble.

If shifts in attention were responsible for the strengthening of DP_{MI} and $DP_{|R|}$ as time approached the motion pulse onset (Fig. 3), then our results are similar to experiments that have shown that correlations decrease with attention. However, previous studies linking attention to reductions in correlation have generally reported positive noise correlation values (Cohen and Maunsell, 2009). Across-trial noise correlations are thought to reflect common sources of feedforward input or connections within the cortical region and feedback from other cortical areas (Ecker et al., 2014). In theory there does not have to be a relationship between noise and single-trial spike-time correlations. Just before the motion pulse, noise correlations in our data were only weakly correlated with average spike-time correlations (Pearson's $R = 0.31$, $p = 0.04$, $N = 47$). Similarly, noise correlations were always positive whether computed using trials with only positive or only negative spike-time correlations (mean noise correlation = 0.12 and 0.09, respectively).

Microsaccades can affect behavioral performance in visually guided tasks (Bair and O'Keefe, 1998; Martinez-Conde et al., 2000) and neural correlations (Snodderly et al., 2001; McFarland et al., 2016). Furthermore, it has been shown that microsaccades may contribute to the link between neural activity and behavior (Herrington et al., 2009). Spike-time correlations in our recordings became stronger with the presence of putative microsaccades, and there were more microsaccades associated with positive versus negative spike-time correlations. However, given that our eye data were filtered at 20 Hz, our conclusions regarding the contribution of microsaccades are limited.

MU activity may lead to overestimations of across-trial correlation values (Cohen and Kohn, 2011). Our single-unit analysis (Fig. 6G,H) showed the same trends as our MU data, but an order of magnitude fewer single-unit spikes produced much noisier DP estimates. Nonetheless, it is difficult to argue that the additional noise contained in MU spike trains would give rise to both positive and negative spike-time correlations that were more predictive of behavior. For example, removing the trials with the largest electrode variances had no appreciable effect on $DP_{|R|}$ values (Fig. 9N). Thus, it is likely that our MU activity mainly reflected the net activity of a pool of similarly tuned neurons.

References

- Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101. [CrossRef Medline](#)
- Averbeck BB, Latham PE, Pouget A (2006) Neuronal correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366. [CrossRef Medline](#)
- Bair W, O'Keefe LP (1998) The influence of fixational eye movements on the response of neurons in area MT of the macaque. *Vis Neurosci* 15:779–786. [Medline](#)
- Bair W, Zohary E, Newsome WT (2001) Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J Neurosci* 21:1676–1697. [CrossRef Medline](#)
- Bosking WH, Maunsell JH (2011) Effects of stimulus direction on the correlation between behavior and single units in MT during a motion detection task. *J Neurosci* 31:8230–8238. [CrossRef Medline](#)
- Bosman CA, Schoffelen JM, Brunet N, Oostenveld R, Bastos AM, Womelsdorf T, Rubehn B, Stieglitz T, De Weerd P, Fries P (2012) Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron* 75:875–888. [CrossRef Medline](#)
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13:87–100. [CrossRef Medline](#)
- Chelaru MI, Dragoi V (2016) Negative correlations in visual cortical networks. *Cereb Cortex* 26:246–256. [CrossRef Medline](#)
- Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN (2011) Variance as a signature of neural computations during decision-making. *Neuron* 69:818–831. [CrossRef Medline](#)
- Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, Bradley DC, Smith MA, Kohn A, Movshon JA, Armstrong KM, Moore T, Chang SW, Snyder LH, Lisberger SG, Priebe NJ, et al (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci* 13:369–378. [CrossRef Medline](#)
- Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nat Neurosci* 14:811–819. [CrossRef Medline](#)
- Cohen MR, Maunsell JH (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12:1594–1600. [CrossRef Medline](#)
- Cohen MR, Maunsell JH (2010) A neuronal population measure of attention predicts behavioral performance on individual trials. *J Neurosci* 30:15241–15253. [CrossRef Medline](#)
- Cohen MR, Maunsell JH (2011) Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron* 70:1192–1204. [CrossRef Medline](#)
- Cohen MR, Newsome WT (2008) Context-dependent changes in functional circuitry in visual area MT. *Neuron* 60:162–173. [CrossRef Medline](#)
- Cook EP, Maunsell JH (2002) Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nat Neurosci* 5:985–994. [CrossRef Medline](#)
- Cook EP, Maunsell JH (2004) Attentional modulation of motion integration of individual neurons in the middle temporal visual area. *J Neurosci* 24:7964–7977. [CrossRef Medline](#)
- Coull JT (2004) fMRI studies of temporal attention: allocating attention within or towards, time. *Cogn Brain Res* 21:216–226. [CrossRef](#)
- de la Rocha J, Doiron B, Shea-Brown E, Josić K, Reyes A (2007) Correlation between neural spike trains increases with firing rate. *Nature* 448:802–806. [CrossRef Medline](#)
- Denfield GH, Ecker AS, Shinn TJ, Bethge M and Tolias AS (2017) Attentional fluctuations induce shared variability in macaque primary visual cortex. *BioRxiv*. Advance online publication. Retrieved April 8, 2018. doi:10.1101/189282. [CrossRef](#)
- deOliveira SC, Thiele A, Hoffmann KP (1997) Synchronization of neuronal activity during stimulus expectation in a direction discrimination task. *J Neurosci* 17:9248–9260. [Medline](#)
- Ditterich J, Mazurek M, Shadlen MN (2003) Microstimulation of visual cortex affects the speed of perceptual decisions. *Nat. Neurosci* 6:891–898. [CrossRef](#)
- Doherty JR, Rao A, Mesulam MM, Nobre AC (2005) Synergistic effect of combined temporal and spatial expectations on visual attention. *J Neurosci* 25:8259–8266. [CrossRef Medline](#)
- Downer JD, Niwa M, Sutter ML (2015) Task engagement selectively modulates neural correlations in primary auditory cortex. *J Neurosci* 35:7565–7574. [CrossRef Medline](#)
- Ecker AS, Berens P, Keliris GA, Bethge M, Logothetis NK, Tolias AS (2010) Decorrelated neuronal firing in cortical microcircuits. *Science* 327:584–587. [CrossRef Medline](#)
- Ecker AS, Berens P, Cotton RJ, Subramaniam M, Denfield GH, Cadwell CR, Smirnakis SM, Bethge M, Tolias AS (2014) State dependence of noise correlations in macaque primary visual cortex. *Neuron* 82:235–248. [CrossRef Medline](#)
- Efron B, Tibshirani R (1986) Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Stat Sci* 1:54–77. [CrossRef](#)
- Farah K, Smith JE, Cook EP (2014) ROC-based estimates of neural-behavioral covariations using matched filters. *Neural Comput* 26:1667–1689. [CrossRef Medline](#)
- Fries P (2009) Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annu Rev Neurosci* 32:209–224. [CrossRef Medline](#)
- Fries P, Reynolds JH, Rorie AE, Desimone R (2001a) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291:1560–1563. [CrossRef Medline](#)
- Fries P, Neuenschwander S, Engel AK, Goebel R, Singer W (2001b) Rapid feature selective neuronal synchronization through correlated latency shifting. *Nat Neurosci* 4:194–200. [CrossRef Medline](#)
- Fries P, Womelsdorf T, Oostenveld R, Desimone R (2008) The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *J Neurosci* 28:4823–4835. [CrossRef Medline](#)
- Gasparini S, Migliore M, Magee JC (2004) On the initiation and propagation of dendritic spikes in CA1 1268 pyramidal neurons. *J Neurosci* 24:11046–11056, 2004. [CrossRef Medline](#)
- Ghose GM, Maunsell JH (2002) Attentional modulation in visual cortex depends on task timing. *Nature* 419:616–620. [CrossRef Medline](#)
- Gold, JJ, Shadlen, MN (2007) The neural basis of decision making. *Annual Review of Neuroscience* 30:535–574. [CrossRef Medline](#)
- Gomez-Ramirez M, Trzcinski NK, Mihalas S, Niebur E, Hsiao SS (2014) Temporal correlation mechanisms and their role in feature selection: a single-unit study in primate somatosensory cortex. *PLoS Biol* 12:e1002004. [CrossRef Medline](#)
- Gu Y, Liu S, Fetsch CR, Yang Y, Fok S, Sunkara A, DeAngelis GC, Angelaki DE (2011) Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron* 71:750–761. [CrossRef Medline](#)
- Gutnisky DA, Dragoi V (2008) Adaptive coding of visual information in neural populations. *Nature* 452:220–224. [CrossRef Medline](#)
- Haefner RM, Gerwinn S, Macke JH, Bethge M (2013) Inferring decoding strategies from choice probabilities in the presence of correlated variability. *Nat Neurosci* 16:235–242. [CrossRef Medline](#)
- Herrero JL, Gieselmann MA, Sanayei M, Thiele A (2013) Attention-induced variance and noise correlation reduction in macaque V1 is mediated by NMDA receptors. *Neuron* 78:729–739. [CrossRef](#)
- Herrington TM, Masse NY, Hachmeh KJ, Smith JE, Assad JA, Cook EP (2009) The effect of microsaccades on the correlation between neural activity and behavior in areas MT, VIP and LIP. *J Neurosci* 29:5793–5805. [CrossRef](#)
- Hirabayashi T, Miyashita Y (2005) Dynamically modulated spike correlation in monkey inferior temporal cortex depending on the feature configuration within a whole object. *J Neurosci* 25:10299–10307. [CrossRef Medline](#)
- Huang X, Lisberger SG (2013) Circuit mechanisms revealed by spike-timing correlations in macaque area MT. *J Neurophysiol* 109:851–866. [CrossRef Medline](#)
- Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J. Neurosci.* 25:10420–10436. [CrossRef](#)
- Jeanne JM, Sharpee TO, Gentner TQ (2013) Associative learning enhances population coding by inverting interneuronal correlation patterns. *Neuron* 78:352–363. [CrossRef Medline](#)
- Katz LN, Yates JL, Pillow JW, Huk AC (2016) Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature* 535:285–288. [CrossRef Medline](#)
- Kimpo RR, Theunissen FE, Doupe AJ (2003) Propagation of correlated activity through multiple stages of a neural circuit. *J Neurosci* 23:5750–5761. [Medline](#)

- Kohn A, Zandvakili A, Smith MA (2009) Correlations and brain states: from electrophysiology to neuroimaging. *Curr Opin Neurobiol* 19:434–438. [CrossRef Medline](#)
- Kohn A, Coen-Cagli R, Kanitscheider I, Pouget A (2016) Correlations and neuronal population information. *Annu Rev Neurosci* 39:237–256. [CrossRef Medline](#)
- Latham PE, Roudi Y (2011) Role of correlations in population coding. In: *Principles of neural coding* (Quiñero R, Panzeri S, eds), pp 121–138. Boca Raton, FL: CRC.
- Lee J, Maunsell JH (2010) The effect of attention on neuronal responses to high and low contrast stimuli. *J Neurophysiol* 104:960–971. [CrossRef](#)
- Losonczy A, Magee JC (2006) Integrative properties of radial oblique dendrites in hippocampal CA1 pyramidal neurons. *Neuron* 50:291–307. [CrossRef Medline](#)
- Martinez-Conde S, Macknik SL, Hubel DH (2000) Microsaccadic eye movements and firing of single cells in the striate cortex of macaque monkeys. *Nat Neurosci* 3:251–258. [CrossRef Medline](#)
- McFarland JM, Cumming BG, Butts DA (2016) Variability and correlations in primary visual cortical neurons driven by fixational eye movements. *J Neurosci* 36:6225–6241. [CrossRef Medline](#)
- Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63:879–888. [CrossRef Medline](#)
- Moon Y, Rajagopalan P, Lall U (1995) Estimation of mutual information using kernel density estimators. *Phys Rev E Stat Nonlin Soft Matter Phys* 52:2318. [CrossRef](#)
- Müller JR, Metha AB, Krauskopf J, Lennie P (2001) Information conveyed by onset transients in responses of striate cortical neurons 21:6978–6990. [CrossRef](#)
- Ni AM, Ruff DA, Alberts JJ, Symmonds J, Cohen MR (2018) Learning and attention reveal a general relationship between population activity and behavior. *Science* 359:463–465. [CrossRef Medline](#)
- Oram MW (2011) Visual stimulation decorrelates neuronal activity. *J Neurophysiol* 105:942–957. [CrossRef Medline](#)
- Palanca BJ, DeAngelis GC (2005) Does neuronal synchrony underlie visual feature grouping? *Neuron* 46:333–346. [CrossRef Medline](#)
- Pang K, Merkel F, Egeth H, Olton DS (1992) Expectancy and stimulus frequency: a comparative analysis in rats and humans. *Percept Psychophys* 51:607–615. [CrossRef Medline](#)
- Parker AJ, Newsome WT (1998) Sense and the single neuron: probing the physiology of perception. *Annu Rev Neurosci* 21:227–277. [CrossRef Medline](#)
- Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci* 5:805–811. [CrossRef Medline](#)
- Pinsker M S (1964) *Information and information stability of random variables and processes*. San Francisco: Holden Day.
- Price NS, Born RT (2010) Timescales of sensory and decision-related activity in the middle temporal and medial superior temporal areas. *J Neurosci* 30:14036–14045. [CrossRef Medline](#)
- Rabinowitz, NC, Goris RL, Cohen M, Simoncelli EP (2015) Attention stabilizes the shared gain of V4 populations. *Elife* 4:e08998. [CrossRef Medline](#)
- Reynolds JH, Pasternak, T, Desimone, R (2000) Attention increases sensitivity of V4 neurons. *Neuron* 26:703–714. [CrossRef](#)
- Ruff DA, Cohen MR (2014) Attention can either increase or decrease spike count correlations in visual cortex. *Nat Neurosci* 17:1591–1597. [CrossRef Medline](#)
- Sato TK, Nauhaus I, Carandini M (2012) Traveling waves in visual cortex. *Neuron* 75:218–229. [CrossRef Medline](#)
- Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16:1486–1510. [CrossRef Medline](#)
- Shuler MG, Bear MF (2006) Reward timing in the primary visual cortex. *Science* 11:1606–1609. [CrossRef](#)
- Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. *Annu Rev Neurosci* 18:555–586. [CrossRef Medline](#)
- Smith JE, Zhan CA, Cook EP (2011) The functional link between area MT neural fluctuations and detection of a brief motion stimulus. *J Neurosci* 31:13458–13468. [CrossRef Medline](#)
- Smith JET, Beliveau V, Schoen A, Remz J, Zhan CA, Cook EP (2015) Dynamics of the functional link between area MT LFPs and motion detection. *J Neurophysiol* 114:80–98. [CrossRef Medline](#)
- Smith MA, Sommer MA (2013) Spatial and temporal scales of neuronal correlation in visual area V4. *J Neurosci* 33:5422–5432. [CrossRef Medline](#)
- Smith PL, Ratcliff R (2004) Psychology and neurobiology of simple decisions. *Trends Neurosci* 27:161–178. [CrossRef Medline](#)
- Snodderly DM, Kagan I, Gur M (2001) Selective activation of visual cortex neurons by fixational eye movements: implications for neural coding. *Vis Neurosci* 18:259–277. [CrossRef Medline](#)
- Sompolinsky H, Yoon H, Kang K, Shamir M (2001) Population coding in neuronal systems with correlated noise. *Phys Rev E Stat Nonlin Soft Matter Phys* 64:051904. [CrossRef Medline](#)
- Steinmetz PN, Roy A, Fitzgerald PJ, Hsiao SS, Johnson KO, Niebur E (2000) Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature* 404:187–190. [CrossRef Medline](#)
- Thiele A, Stoner G (2003) Neuronal synchrony does not correlate with motion coherence in cortical area MT. *Nature* 421:366–370. [CrossRef Medline](#)
- Wimmer K, Compte A, Roxin A, Peixoto D, Renart A, de la Rocha J (2015) Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT. *Nat Commun* 6:6177. [CrossRef Medline](#)
- Womelsdorf T, Fries P (2007) The role of neuronal synchronization in selective attention. *Curr Opin Neurobiol* 17:154–160. [CrossRef Medline](#)
- Wright BA, Fitzgerald MB (2004) The time course of attention in a simple auditory detection task. *Percept Psychophys* 66:508–516. [CrossRef Medline](#)
- Zandvakili A, Kohn A (2015) Coordinated neuronal activity enhances corticocortical communication. *Neuron* 87:827–839. [CrossRef Medline](#)
- Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–143. [CrossRef Medline](#)
- Zylberberg A, Fetsch CR, Shadlen MN (2016) The influence of evidence volatility on choice, reaction time and confidence in a perceptual decision. *Elife* 5:e17688. [CrossRef Medline](#)