Simultaneous multi-area recordings suggest a novel hypothesis about how attention improves performance

Douglas A. Ruff and Marlene R. Cohen

Department of Neuroscience and Center for the Neural Basis of Cognition, University of Pittsburgh, Pittsburgh, Pennsylvania 15260; email: ruffd@pitt.edu

Abstract - Visual attention dramatically improves subjects' ability to see and also modulates the responses of visual and oculomotor neurons. Despite hundreds of studies demonstrating the co-occurrence of behavioral and neuronal effects of attention [1], the relationship between neuronal modulations and improved performance remains unknown. There are three dominant hypotheses (Figure 1A): that attention 1) improves visual information coding [2,3], 2) improves communication between visual and oculomotor brain areas [4], or 3) reduces uncontrolled fluctuations in the animals' cognitive state, thereby improving performance and rendering observed effects of attention on response variability as epiphenomena [5]. Here we show that none of these hypotheses account for observed perceptual improvements and we provide evidence for a novel hypothesis: that the well known effects of attention on firing rates and shared response variability in visual cortex reshape the representation of attended stimuli such that they more effectively drive downstream neurons and guide decisions without explicitly changing the weights relating sensory responses to downstream neurons or behavior (Figure 1B). We assessed each hypothesized mechanism by recording simultaneously from groups of neurons in two stages of visuomotor processing: area MT, which encodes motion information [6] and the superior colliculus (SC), where neuronal responses are either visual, oculomotor, or intermediate [7-9] and have been hypothesized to be involved in computing perceptual decisions [10]. Constraining our analyses by the animals' behavior and the simultaneous recordings from both areas allowed us to reject each prior hypothesis and greatly clarified the relationship between attention, neuronal responses and behavior. Our results suggest a path toward understanding the neural underpinnings of

perception and cognition in health and disease by analyzing neuronal responses in ways that are constrained by behavior and interactions between brain areas.

Main - We compared evidence for and against three hypothesized attention mechanisms using neuronal responses collected while two rhesus monkeys performed a widely studied motion direction change-detection task Figure 1C; [11,12,13]. The animals' performance was greatly affected (Figure 1D) by a cue instructing them to shift spatial attention between a stimulus within the same or opposite hemifield as the joint receptive fields of several dozen neurons that were recorded on multielectrode probes in MT (Figure 1E, red points) and the SC (blue points). MT and the SC represent different stages of perceptual decision-making and therefore provide the opportunity to evaluate each hypothesized attention mechanism. MT contributes to motion perception [6,14]. The SC is thought to play many roles in visually guided tasks [7-9], decision-making [10,15] and attention[16].

Figure 1

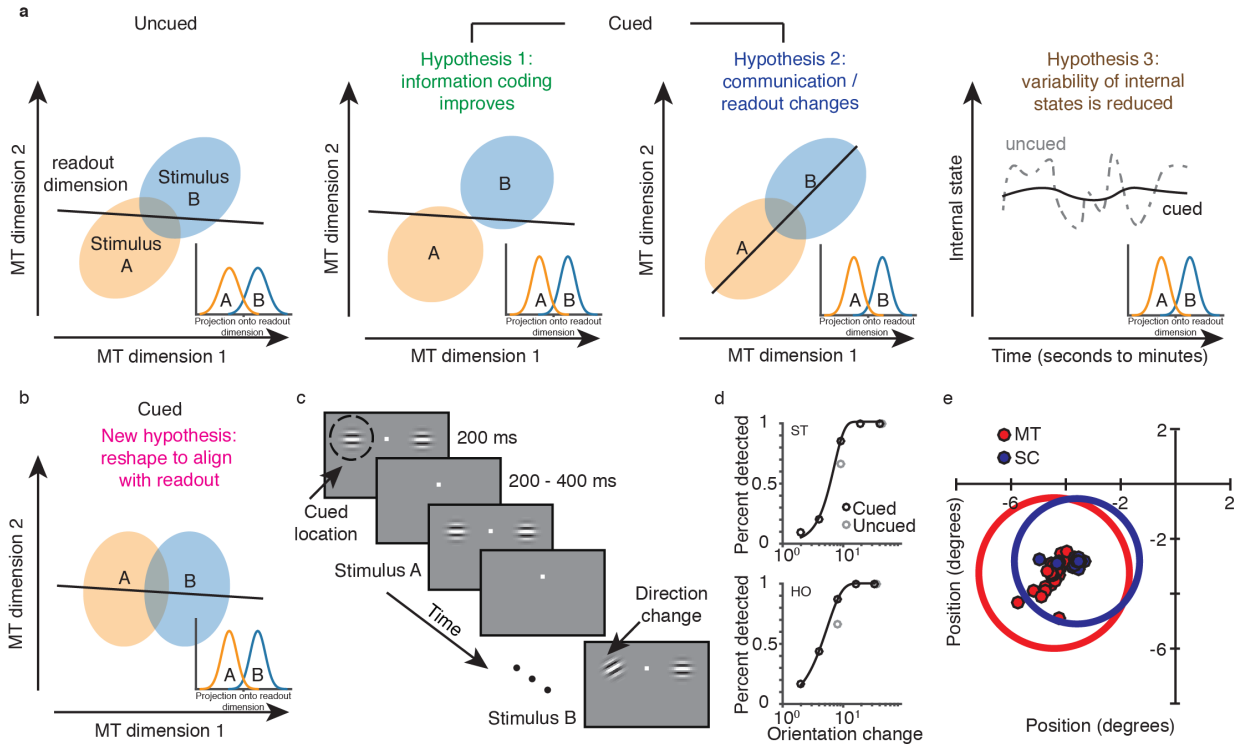

Figure 1. Hypotheses and methods. (A) Schematics describing predominant hypotheses about links between attention, visual cortical activity, and behavior. The left plot depicts MT population responses to two visual stimuli plotted along two dimensions in population response space (e.g. the first two principal components; see Methods) and a readout dimension which represents the visual information that is communicated to neuronal populations involved in planning behavior during the uncued condition. The insets depict projections of the population responses onto the readout dimension. Hypothesis 1 is that the MT representations of the two stimuli become more easily distinguishable (e.g. by separating the distributions of responses to the two stimuli). In this scenario, the distributions of projections along even a suboptimal readout axis may also be more separable. Hypothesis 2 suggests that attention changes the communication or readout dimensions, such that projections of responses to the two stimuli are more separable. Hypothesis 3 proposes that attention-related changes in visual cortex are epiphenomenal. Instead, it suggests that attention reduces uncontrolled variability in the animals' internal states, which might produce less variable neuronal population responses and therefore more separable projections onto the readout dimensions. (B) Our new hypothesis: attention reshapes population responses so they are better aligned with relatively static readout dimensions. (C) Direction change-detection task with cued attention. The stimuli before the change can be thought of as stimulus A and the changed stimulus can be thought of as stimulus B in the schematics in (A). (D) Psychometric curves from two example sessions (monkey ST, top, monkey HO, bottom) with best-fitting Weibull functions. Attention improved detection of median difficulty trials by 25% on average across all experiments. (E) Receptive field (RF) centers of recorded units from the same example session as in the top plot in (D). Dots represent the RF center (red, MT; blue, SC). The circle represents the size and location of the median RF from each area.

The three predominant attention hypotheses make different predictions about how attention should affect MT and the SC in our task. The first (information coding) hypothesis predicts that attention improves the motion direction information encoded in MT. The second (communication) hypothesis posits that attention improves the stimulus information that is shared between MT and the SC. The third (internal states) hypothesis suggests that attention reduces uncontrolled fluctuations in the animals' cognitive states, allowing them to focus on their psychophysical task. In this scenario, the widely observed attention-related reduction in response variability in visual cortex [11,13,17] would be an epiphenomenon of decreased variability in cognitive states.

Past studies have evaluated these hypotheses by analyzing the responses of individual neurons or pairs of neurons, which typically lack the statistical power to reveal a strong link to behavior. Using our data set, we replicated the observations that have been used as evidence in favor of all three hypotheses. Consistent with previous studies evaluating the information coding hypothesis [2,3,18], we found that attention increased the trial-averaged responses of neurons in both MT and the SC (Supplemental Figure 1A and B). Consistent with studies evaluating both the information coding and internal states hypotheses [5], we found that attention decreased the extent to which the trial to trial fluctuations in neuronal responses to repeated presentations of the same stimulus are shared between pairs of MT neurons[11,13,17] (quantified as the average spike count or noise correlation, or r_{SC}[19]; Supplementary Figure 1C). Consistent with studies evaluating the communication hypothesis, attention increases correlated variability between the two areas [12,20] (Supplemental Figure 1C). This attention-related increase was weakly dependent on the visual responsivity of SC neurons (Supplemental Figure 2).

Most previous studies aim to connect the observed attention-related changes in single neurons or pairs of neurons to behavior via an explicit or implicit model. Recordings from neuronal populations provide the opportunity to do so directly. We reasoned that analyzing the relationship between populations of simultaneously recorded neurons in multiple brain areas with the animals' behavior would provide insight into the relative importance of each hypothesized mechanism.

The internal states hypothesis provides a challenge for the approach of linking populations of visual neurons with behavior because it proposes that there is no link between attention-related changes in MT or other visual areas and performance. Instead, it posits that attention limits mind wandering, and the changes in performance and neuronal responses simply reflect improved stability in internal states [5]. There are two reasons this hypothesis seems unlikely to account for our data. First, the effects of spatial attention are spatially specific (e.g. correlated variability increases in one hemisphere while decreasing in the other, even when neurons in the two hemispheres are simultaneously recorded[11]), meaning that reductions in the variability of global cognitive processes like arousal and motivation are unlikely to account for the attention-related changes in visual cortex. Further, it is not obvious how reductions in fluctuations in internal states could account for the attention-related increases in firing rates observed in spatial attention studies like ours (Supplementary Figure 1), let alone the more complex firing rate changes associated with feature attention [21].

This hypothesis can also be addressed using a population-analysis approach, by using the responses of MT and SC neurons to attempt to quantify the variability in the animals' internal states. We reasoned that fluctuations in internal states would 1) often occur at timescales longer than the 400-600 ms between stimulus presentations in our task and 2) affect the covariability of neurons in many brain areas, including both MT and the SC. We performed principal components analysis on population responses to the identical visual stimuli that occurred before the direction change on each trial (e.g. stimulus A in Figure 1; see Methods) in each attention condition. Because the only variability in those population responses is internally generated, the first principal component (PC) represents the axis of greatest shared variability in the population of neurons in each brain area.

Figure 2

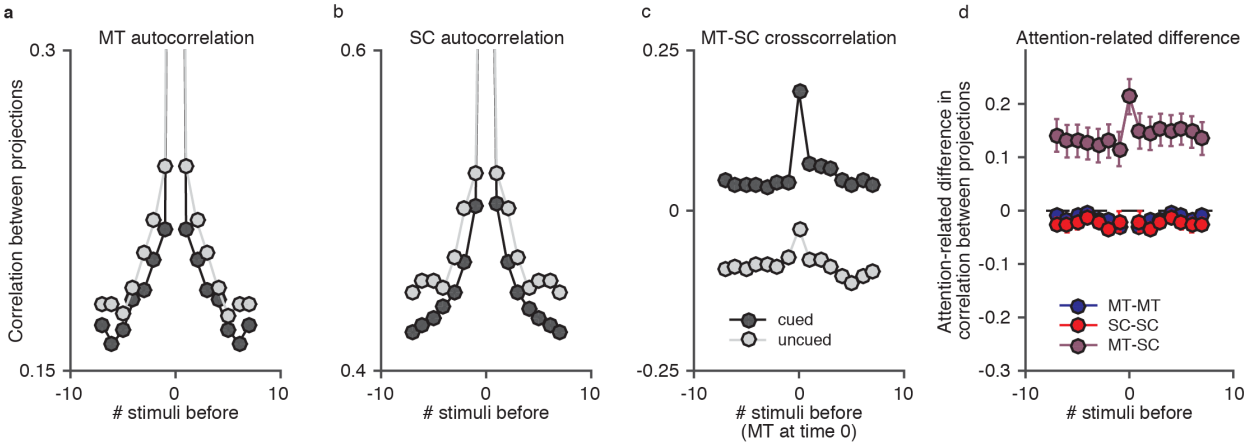

Figure 2. Attention has opposite effects on slow fluctuations in neuronal population responses within and across areas. (A, B) Autocorrelations between projections onto the first PCs of population responses to repeated presentations of the same visual stimulus in (A) MT, and (B) the SC. The x-axis plots time lag in units of stimulus presentations (400-600 ms; see Methods). (C) Cross correlation between projections onto the first PCs in MT and the SC (same data and plotting conventions as in A and B). (D) Attention-related difference in autocorrelation or cross correlations between the projections in the previous plots. Error bars represent standard error of the mean. Attention was associated with a statistically significant decrease in autocorrelation overall (t-tests, p<.05) in both areas and in 11/15 individual MT data sets and 9/15 SC data sets (t-tests, p<0.05 with a Bonferroni correction) and a significant increase in cross correlation overall (t-test, p<.001) and in 11/15 individual data sets.

The autocorrelation functions of projections onto the first PC show that there is indeed response variability in each area that fluctuates slowly and is reduced by attention (Figures 2A, B), which is not as readily observable by computing noise correlations between pairs of neurons (Supplementary Figure 1C). However, the cross-correlation of projections onto the first PC in MT and the SC showed a qualitatively different time course than the autocorrelation function. Furthermore, not only did attention not reduce the covariability of these signals between these two areas (and presumably brain wide), attention *increased* the slow variability that is shared between areas (Figure 2C and Supplementary Figure 1C). These results are in conflict with the idea that the attention-related decrease in covariability within each area is a byproduct of a decrease in uncontrolled fluctuations in internal states, because such a decrease should be brain-wide.

Figure 3

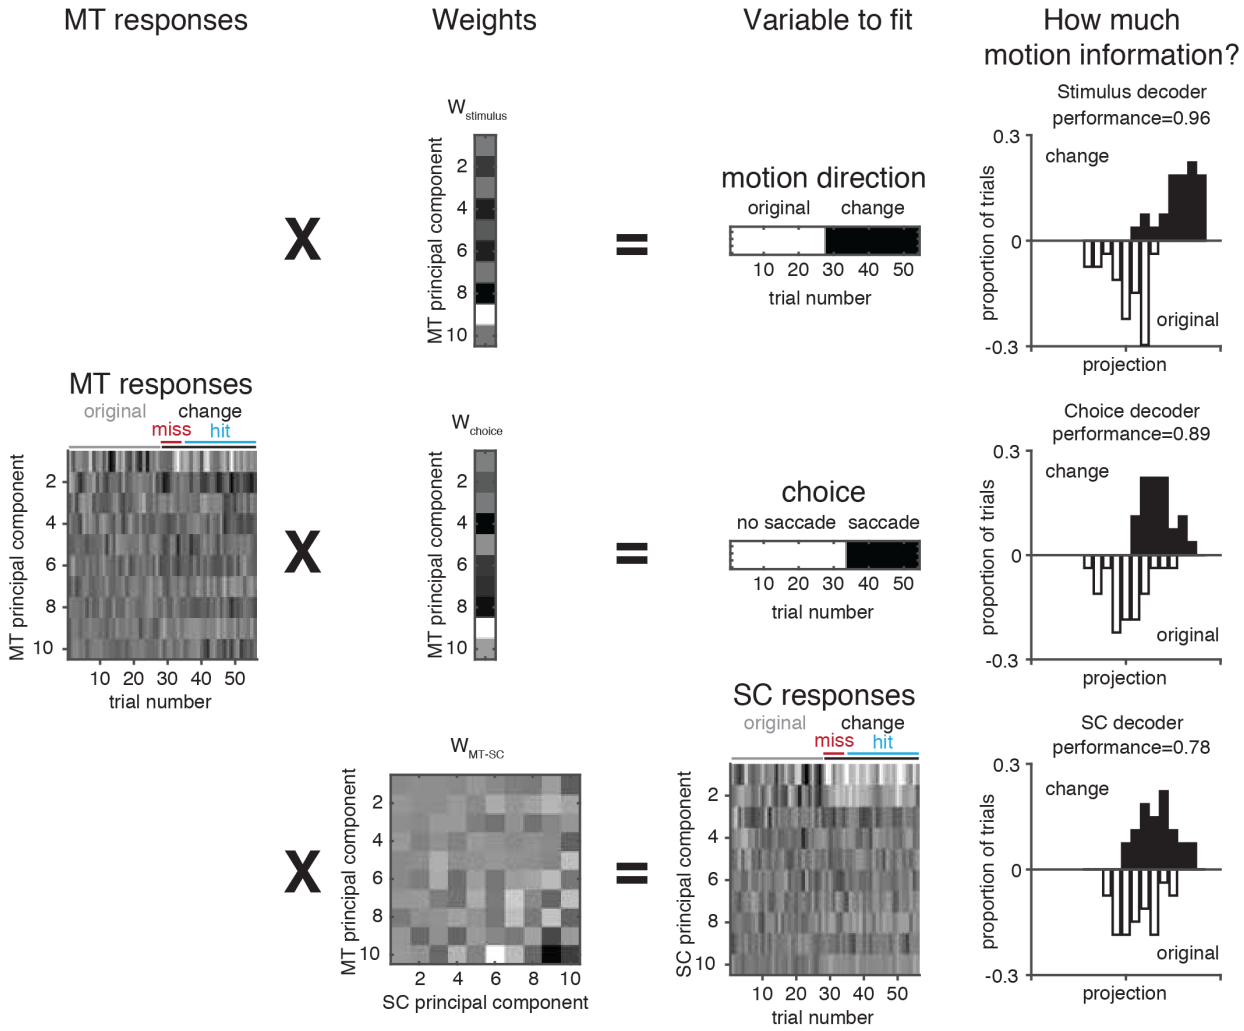

Figure 3 – Schematic of our decoding procedure. We used linear regression to find the weights (second column) that best relate the ***first ten principal components of the MT population's response*** *(left) to the visual stimulus (Stimulus decoder; top row), the animal's choice (Choice decoder; middle row), or the projections of the responses of the population of simultaneously recorded SC neurons (SC decoder; bottom row). We assessed the performance of each decoder by decoding stimulus information from MT responses on a separate set of trials using each set of weights (right column). See methods for detailed decoding and cross validation procedures.*

To investigate the remaining hypotheses, we determined whether attention affects the amount of stimulus information that can be decoded from the population of MT neurons using linear decoders that are optimized to a) dissociate between the original and changed stimuli (Stimulus decoder in Figure 3), b) predict the animals' choices (whether or not they made an eye

movement; Choice decoder), or c) predict the activity of the population of SC neurons we recorded (SC decoder).

The information coding hypothesis posits that attention improves the stimulus information that could be gleaned by an optimal stimulus decoder. However, attention did not significantly affect the performance of an optimal decoder in our data set, even when we used a decoder optimized separately for each attention condition (Figure 4A, left bars). Furthermore, theoretical work suggests that the effects of attention on the stimulus information that can be decoded from small neuronal populations are likely to be even more minimal for larger populations [22-24].

The communication hypothesis posits that attention improves communication between visual and decision areas, meaning that it should change the weights relating MT responses to either behavior or SC responses. We found that attention had larger effects on the stimulus information that is related to the animals' choices on individual trials (Figure 4A, middle bars) or that is shared with the SC (Figure 4A, right bars). However, this could arise from either a weight change (Figure 1A) or a change within MT, to better align MT responses with static readout weights (Figure 1B).

Figure 4. Effects of attention on the stimulus information that can be decoded from small populations of MT neurons. (A) Ability of a cross-validated linear decoder to distinguish the original from changed stimuli for each decoder. Error bars represent SEM. The effect of attention was significant for the Choice and SC decoders (paired t-tests, p<.05) but not for the Stimulus decoder (p=0.28). The effects of attention on the Choice and SC decoders were greater than for the stimulus decoder (paired t-tests, p<.05), but not significantly different from each other (p=0.21). (B) Weight swapping analysis demonstrates that decoding performance was typically better using the MT responses from the cued condition and the choice decoder weights from the uncued condition (y-axis) than using the MT responses from the uncued condition and the Choice decoder weights from the cued condition (x-axis; paired t-test, p<.05). (C) Same, using the weights from the SC decoder (paired t-test, p<.05).

Figure 4

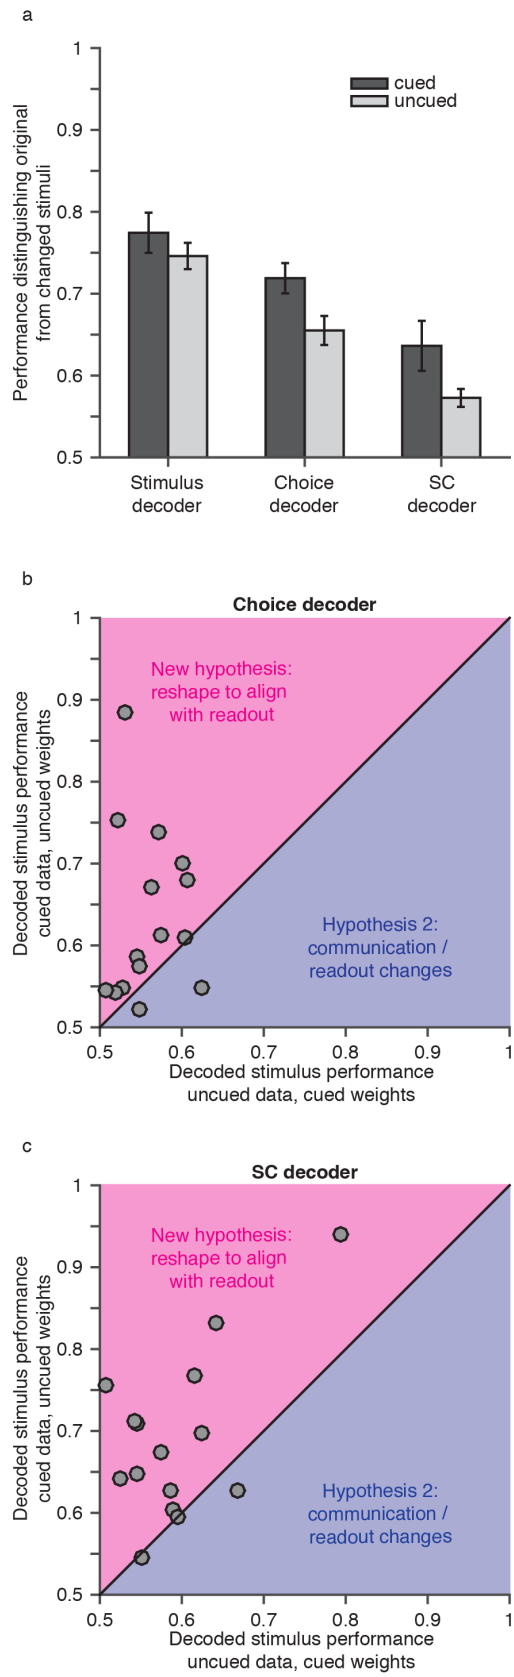

We found no evidence that attention changes weights relating MT responses to SC responses or behavior. Both the Choice and SC decoders gleaned more stimulus information from MT responses in the attended than unattended condition, even when we used the weights computed in the opposite attention condition for which they were calculated (Figures 4B and 4C). Together, these neuronal population analyses that use the animals' behavior and the activity of downstream neurons to assess the hypothesized attention mechanisms reveal that none of the three previously hypothesized mechanisms provides a

satisfactory account of the observed attention-related behavioral improvement.

Our data support the hypothesis that attention reshapes the representation of attended stimuli to more effectively guide behavior (Figure 1B). In this scenario, the critical changes are in visual cortex. However, this does not result in a large improvement in the stimulus information that can be gleaned by an optimal stimulus decoder. Instead, the modulated neuronal activity in MT better aligns with the readout dimensions using relatively static weights. Modifying stimulus representations without changing readout dimensions may be simple mechanistically[25], allowing animals to discriminate stimulus features (e.g. motion direction) of any stimuli they encounter in the natural world, without optimizing readout for the reduced stimulus sets used in laboratory tasks like ours [18].

It is important to note that not only is this reshaping of the representation of the attended stimulus consistent with the simple signatures of attention that have been observed for many years (e.g. changes in firing rate gain or noise correlations [2]), it may be a direct result of those changes. For example, we showed that in this task, monkeys' choices are well-aligned with the axis in population space that explains the most correlated noise [13], perhaps because that axis is well-positioned to decode the motion direction of any stimulus [18]. If so, reducing noise correlations and increasing firing rate gains would improve the stimulus information projected along that readout axis (following the intuitions in [26]). Attention-related changes in gains and correlations have been observed for many years. The key insight that they may improve performance by aligning the responses of populations of visual neurons to fixed readout

235 dimensions was only possible using population analyses constrained by the animals' behavior
236 rather than focusing on one or two neurons.
237
238 By combining psychophysics and multi-neuron, multi-area recordings, our study provides a
239 framework for understanding the neuronal mechanisms underlying a wide variety of sensory,
240 cognitive, and motor processes. We showed that the effects of covert spatial attention are best
241 described as a change in the visual information used to guide behavior. In the future, it will be
242 interesting to use the same approach to determine whether a similar mechanism can account for
243 behavioral changes associated with other cognitive processes (e.g. feature attention) that might
244 seem more likely to change the weights relating stimulus information to downstream neurons or
245 behavior. Further, many neuropsychiatric disorders (including disorders of attention, Autism, and
246 schizophrenia) are thought to involve changes in the same computations thought to underlie
247 attention [27]. An exciting possibility is that these changes might be identified and potential
248 therapies evaluated in animal models using the combination of behavioral evaluation and multi-
249 neuron, multi-area recordings that we described here.
250
251 Methods Summary
252
253 The subjects were two adult male rhesus monkeys (*Macaca mulatta*, 8 and 9 kg). All animal
254 procedures were approved by the Institutional Animal Care and Use Committees of the
255 University of Pittsburgh and Carnegie Mellon University. Using linear 24 channel moveable
256 probes (Plexon), we simultaneously recorded extracellular activity from direction-selective
257 neurons in area MT and neurons in the superior colliculus that responded either visually, prior to

258 a saccade, or both. Before beginning the experiment, we searched for neurons in both areas that
259 had overlapping spatial receptive fields (Figure 1E) as determined by mapping with both drifting
260 gratings and a delayed saccade task. The monkeys performed a direction change-detection task
261 that commenced upon fixation of a central spot (Figure 1C). Two drifting Gabor stimuli, whose
262 direction was selected to drive the recorded population of MT neurons well, flashed on and off
263 until the direction of one stimulus changed at a random, unsignaled time. The monkeys signaled
264 detection of the change by making a saccade to the changed stimulus within 450 ms of its onset.
265 The location of the stimulus change was cued using instruction trials prior to each block of trials
266 and the cue was valid 80% of the time. On the other 20% of trials, the change happened at the
267 uncued location. In order to encourage fixation on longer trials, catch trials, where no stimulus
268 changed direction and monkeys were rewarded for maintaining fixation, were intermixed. We
269 analyzed spiking activity in response to all visual stimuli except the first stimulus in each trial.

270

271 Materials and Methods

272 The subjects were two adult male rhesus monkeys (*Macaca mulatta*, 8 and 9 kg). All animal
273 procedures were approved by the Institutional Animal Care and Use Committees of the
274 University of Pittsburgh and Carnegie Mellon University.

275 We presented visual stimuli using custom software (written in MATLAB using the
276 Psychophysics Toolbox [28,29] on a CRT monitor (calibrated to
277 linearize intensity; 1024 768 pixels; 120 Hz refresh rate) placed 54 cm from the animal. We
278 monitored eye position using an infrared eye tracker (Eyelink 1000; SR Research) and recorded
279 eye position and pupil diameter (1000 samples/s), neuronal responses (30,000 samples/s), and

280 the signal from a photodiode to align neuronal responses to stimulus presentation times (30,000
281 samples/s) using hardware from Ripple.
282
283 Behavioral Task

284 As previously described [11], a trial began when the monkey fixated a small, central spot within a
285 1.25° per side, square fixation window in the center of a video display while two peripheral
286 Gabor stimuli (one overlapping the receptive fields of the recorded neurons, the other in the
287 opposite visual hemifield) synchronously flashed on (for 200 ms) and off (for a randomized
288 period between 200-400 ms) until, at a random, unsignaled time, the direction of one of the
289 stimuli changed from that of the preceding stimuli (Figure 1C). The monkey received a liquid
290 reward for making a saccade to the stimulus that changed. Attention was cued in blocks of 50-
291 100 trials, and alternated between blocks where attention was cued to either the left or the right
292 stimulus. In each block, the direction change occurred at the cued stimulus on 80% of trials, and
293 at the uncued stimulus in 20% of trials (all uncued changes used either the middle or largest
294 orientation change, Figure 1D). Catch trials, where no stimulus changed direction and the
295 monkey was rewarded for maintaining fixation, were randomly intermixed throughout each
296 block. Psychometric data were fit with Weibull functions. Before recording commenced, the
297 monkeys were extensively trained to have stable thresholds across a range of spatial locations (3-
298 6 months). During recording, the size, location, spatial frequency and direction of the Gabor
299 stimuli were selected to drive the recorded units well and the direction of all of the stimuli prior
300 to the orientation change was constant throughout a recording session.
301
302 Electrophysiological Recordings

We simultaneously recorded extracellularly from single units and sorted multiunit clusters (the term "unit" refers to either). The dataset consisted of a total of 306 responsive MT units and 345 responsive SC units total (36-58 units per session, mean 20 in MT, 24 in the SC for Monkey HO; 36-53 units per session, mean 21 in MT, 22 in SC for Monkey ST) in both MT and the SC in the right hemisphere using moveable, linear 24-channel V-probes (Plexon; inter-electrode spacing in MT = 50μm, SC = 100μm). We presented visual stimuli and tracked eye position as previously described [12]. The data presented are from 6 days of recording for Monkey HO and 9 days of recording for Monkey ST. Each day consisted of multiple blocks of the attention task (Figure 1C; mean 1015 of trials for Monkey HO, 745 for Monkey ST) preceded by receptive field mapping using a delayed saccade task and direction tuning during passive fixation.

Data Analysis

All spike sorting was done offline manually using Offline Sorter (version 3.3.5; Plexon). We based our analyses on both single units and multiunit clusters and use the term "unit" to refer to either. All neuronal analyses prior to Figure 3 used spike count responses between 50-250 ms after stimulus onset to account for visual latencies in the two areas. To remove response contamination from eye movements during change stimuli, data presented in the decoding analyses in Figure 3 and 4 used shorter response windows. Responses to both unchanged and changed stimuli were measured from 50-185 ms after stimulus onset for monkey HO and 50-220 ms for monkey ST. These times were selected based on the distribution of each animal's reaction times, with the goal of maximizing the number of trials that could be included in the analyses. Trials with reaction times that began during those windows were excluded. To minimize the impact of adaptation on our results, we did not analyze the first stimulus presentation in each

trial. We only analyzed a recorded MT unit if its stimulus-driven firing rate was 10% higher than its firing rate as measured in the 100 ms prior to the onset of the first stimulus. We only analyzed a recorded SC unit if its stimulus-driven firing rate was 10% higher than its firing rate as measured in the 100 ms prior to the onset of the first stimulus or if its response during a 100 ms epoch prior to a saccade on correct trials to the contralateral side was 10% larger than that same baseline. Stimulus presentations during which a microsaccade was detected were excluded from analyses [12,30]).

For firing rate analyses in Supplemental Figure 1A and B, attention indices were calculated using average spike counts on the stimulus presentation prior to correct detections of the intermediate change amount depending on whether attention was directed into or out of the receptive fields of the recorded neurons using the formula $(attend_{in} - attend_{out})/(attend_{in} + attend_{out})$. Significance of individual units was determined by a paired t-test ($p<0.05$).

Noise correlations

We defined the correlated variability of each pair of simultaneously recorded units (quantified as spike count correlation or r_{SC} [19]) as the Pearson correlation coefficient between the responses of the two units to repeated presentations of the same stimulus. This measure of r_{SC} represents noise correlations rather than signal correlations because the responses used in this analysis were always to an identical visual stimulus. For Supplemental Figure 1C, we included responses from stimulus presentations 2 though 10 from trials that ended with either a hit, miss or correct catch trial and that were immediately followed by the maintenance of fixation and continuation of the

trial. We z-scored responses as a function of the stimulus presentation number in each trial and then pooled data across stimulus presentations before calculating noise correlations. Results did not qualitatively change if we did not perform this z-score procedure. For Supplemental Figure 1D, we included data from all stimulus presentations prior to the change stimulus (except the first) and sorted them depending on what the behavioral outcome was on the subsequent stimulus presentation. Pairs of units that were recorded on the same electrode were not included in correlation analyses. The data presented in Supplementary Figures 1C consisted of 3,285 MT pairs, 3,948 SC pairs and 6,934 between area pairs.

Slow timescale auto- and cross-correlations

To test the hypothesis that attention affects uncontrolled fluctuations in internal states, we created a procedure to identify slow fluctuations in population responses. We plotted the responses of the populations of simultaneously recorded MT or SC neurons to the stimuli before the motion direction change (excluding the first stimulus) in a high dimensional space in which each neuron's response comprises one axis. We performed principal components analysis (PCA) on this cloud of points. Because the visual stimuli in this subset of the data were all identical, the first principal component (PC) represents the axis that captures the most variance in non-stimulus related population responses. We searched for slow fluctuations in internal states by measuring the auto- and cross-correlations in projections onto this first PC in each area.

Decoding

We focused our decoding analyses on trials in which the third largest (middle) direction change occurred, because changes of that magnitude occurred in both attention conditions. We used the decoding strategy schematized in Figure 3. We began by constructing four matrices in each attention condition: 'MT responses' (a # MT neurons x 2*# trials matrix of MT responses to the stimuli before the direction change and the changed stimulus on the relevant trials), 'motion direction' (a 1 x 2*# trials vector of zeros for the stimulus before the change, referred to as 'original', and ones for the changed stimulus, referred to as 'change'), 'choice' (a 1 x 2*# trials vector of zeros for stimulus presentations on which the animal did not make an eye movement, referred to as 'no saccade', and ones when the animal made an eye movement, referred to as 'saccade'), and 'SC responses' (a # SC neurons x 2*# trials matrix of SC responses to the stimuli before the direction change and the changed stimulus on the relevant trials).

We cross validated by holding out the two stimulus presentations (for the original and changed stimuli) from one trial at a time to perform the rest of our analyses. To reduce the number of weights we needed to fit and therefore improve our confidence in the weights we did fit, we performed PCA on the MT and SC responses to find the first 10 PCs in each area. We then performed linear regression to find the weight vectors (for the Stimulus and Choice decoders) or weight matrices (for the SC decoder) that related projections along the first ten MT PCs plus a vector of ones to 'motion direction', 'choice', or projections along the first 10 SC PCs in each attention condition.

We assessed the stimulus information in each decoder (Figure 4) by multiplying projections of MT responses to the original and changed stimuli from the held-out trial by the fitted weights

and either determining whether those weighted sums correctly classified the stimuli as original or changed (Stimulus and Choice decoders) or whether a linear classifier correctly classified those stimulus presentations on the basis of the predicted SC responses (SC decoder).

Data Availability

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

Code Availability

The MATLAB code that was written to perform the analyses in this paper are available from the corresponding author upon reasonable request.

References:

1 Carrasco, M. Visual attention: the past 25 years. *Vision Res* **51**, 1484-1525, doi:10.1016/j.visres.2011.04.012 (2011).

2 Maunsell, J. H. R. Neuronal Mechanisms of Visual Attention. *Annu Rev Vis Sci* **1**, 373-391, doi:10.1146/annurev-vision-082114-035431 (2015).

3 Reynolds, J. H. & Chelazzi, L. Attentional modulation of visual processing. *Annual Review of Neuroscience* **27**, 611-647 (2004).

4 Fries, P. Rhythms for Cognition: Communication through Coherence. *Neuron* **88**, 220-235, doi:10.1016/j.neuron.2015.09.034 (2015).

5 Ecker, A. S., Denfield, G. H., Bethge, M. & Tolias, A. S. On the Structure of Neuronal Population Activity under Fluctuations in Attentional State. *J Neurosci* **36**, 1775-1789, doi:10.1523/JNEUROSCI.2044-15.2016 (2016).

6 Born, R. T. & Bradley, D. C. Structure and function of visual area MT. *Annual review of neuroscience* **28**, 157-189, doi:10.1146/annurev.neuro.26.041002.131052 (2005).

7 Wurtz, R. H. & Albano, J. E. Visual-motor function of the primate superior colliculus. *Annual review of neuroscience* **3**, 189-226, doi:10.1146/annurev.ne.03.030180.001201 (1980).

8 Sparks, D. L. Conceptual issues related to the role of the superior colliculus in the control of gaze. *Current opinion in neurobiology* **9**, 698-707, doi:10.1016/S0959-4388(99)00039-2 (1999).

9 Gandhi, N. J. & Katnani, H. A. motor functions of the superior colliculus. *Annual Review of Neuroscience* **34**, 205-231, doi:10.1016/j.micinf.2011.07.011.Innate (2011).

10 Horwitz, G. D. & Newsome, W. T. Separate Signals for Target Selection and Movement Specification in the Superior Colliculus. *Science* **284**, 1158-1161, doi:10.1126/science.284.5417.1158 (1999).

11 Cohen, M. R. & Maunsell, J. H. R. Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience* **12**, 1594-1600 (2009).

12 Ruff, D. A. & Cohen, M. R. Attention Increases Spike Count Correlations between Visual Cortical Areas. *Journal of Neuroscience* **36**, 7523-7534, doi:10.1523/JNEUROSCI.0610-16.2016 (2016).

13 Ni, A. M., Ruff, D. A., Alberts, J. J., Symmonds, J. & Cohen, M. R. Learning and attention reveal a general relationship between population activity and behavior. *Science* **359**, 463-465 (2018).

14 Parker, A. J. & Newsome, W. T. Sense and the single neuron: probing the physiology of perception. *Annual review of neuroscience* **21**, 227-277, doi:10.1146/annurev.neuro.21.1.227 (1998).

15 Horwitz, G. D. & Newsome, W. T. Target selection for saccadic eye movements: prelude activity in the superior colliculus during a direction-discrimination task. *Journal of neurophysiology* **86**, 2543-2558 (2001).

16 Krauzlis, R. J., Lovejoy, L. P. & Zénon, A. Superior Colliculus and Visual Spatial Attention. *Annual Review of Neuroscience* **36**, 165-182, doi:10.1146/annurev-neuro-062012-170249.Superior (2013).

17 Mitchell, J. F., Sundberg, K. A. & Reynolds, J. H. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* **63**, 879-888 (2009).

18 Ruff, D. A., Ni, A. M. & Cohen, M. R. Cognition as a Window into Neuronal Population Space. *Annu Rev Neurosci* **41**, 77-97 (2018).
19 Cohen, M. R. & Kohn, A. Measuring and interpreting neuronal correlations. *Nature Neuroscience* **14**, 811-819 (2011).
20 Ruff, D. A. & Cohen, M. R. A normalization model suggests that attention changes the weighting of inputs between visual areas. *Proc Natl Acad Sci U S A* **114**, E4085-E4094, doi:10.1073/pnas.1619857114 (2017).
21 Treue, S. & Martinez-Trujillo, J. C. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* **399**, 575-579 (1999).
22 Moreno-Bote, R. *et al.* Information-limiting correlations. *Nature Neuroscience* **17**, 1410-1417 (2014).
23 Kanitscheider, I., Coen-Cagli, R. & Pouget, A. Origin of information-limiting noise correlations. *Proceedings of the National Academy of Sciences of the United States of America* **112**, E6973-6982, doi:10.1073/pnas.1508738112 (2015).
24 Kohn, A., Coen-cagli, R., Kanitscheider, I. & Pouget, A. Correlations and Neuronal Population Information. *Annual Review of Neuroscience* **39**, 237-256, doi:10.1146/annurev-neuro-070815-013851 (2016).
25 Huang, C. *et al.* Circuit models of low dimensional shared variability in cortical networks. *bioRxiv* **217976**, doi:https://doi.org/10.1101/217976 (2017).
26 Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nature Reviews Neuroscience* **7**, 358-366 (2006).
27 Carandini, M. & Heeger, D. J. Normalization as a canonical neural computation. *Nature Reviews Neuroscience* **13**, 51-62, doi:10.1038/nrn3136 (2012).
28 Brainard, D. H. The Psychophysics Toolbox. *Spatial vision* **10**, 433-436 (1997).
29 Pelli, D. The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial vision* **10**, 437-442 (1997).
30 Engbert, R. & Kliegl, R. Microsaccades uncover the orientation of covert attention. *Vision Research* **43**, 1035-1045, doi:10.1016/S0042-6989(03)00084-1 (2003).

Acknowledgements

We are grateful to Karen Mckracken for providing technical assistance as well as Faisal Baqai, Matt Getz, Jay Hennig, Chengcheng Huang, Amy Ni and John Maunsell who provided comments on an earlier version of this manuscript.

Author Contributions

Both authors conceived and designed the experiments, analyzed the data and wrote the manuscript. D.A.R. collected the data.

Author Information

Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing interests. Correspondence and requests for materials should be addressed to ruffd@pitt.edu.

Supplemental Figure 1

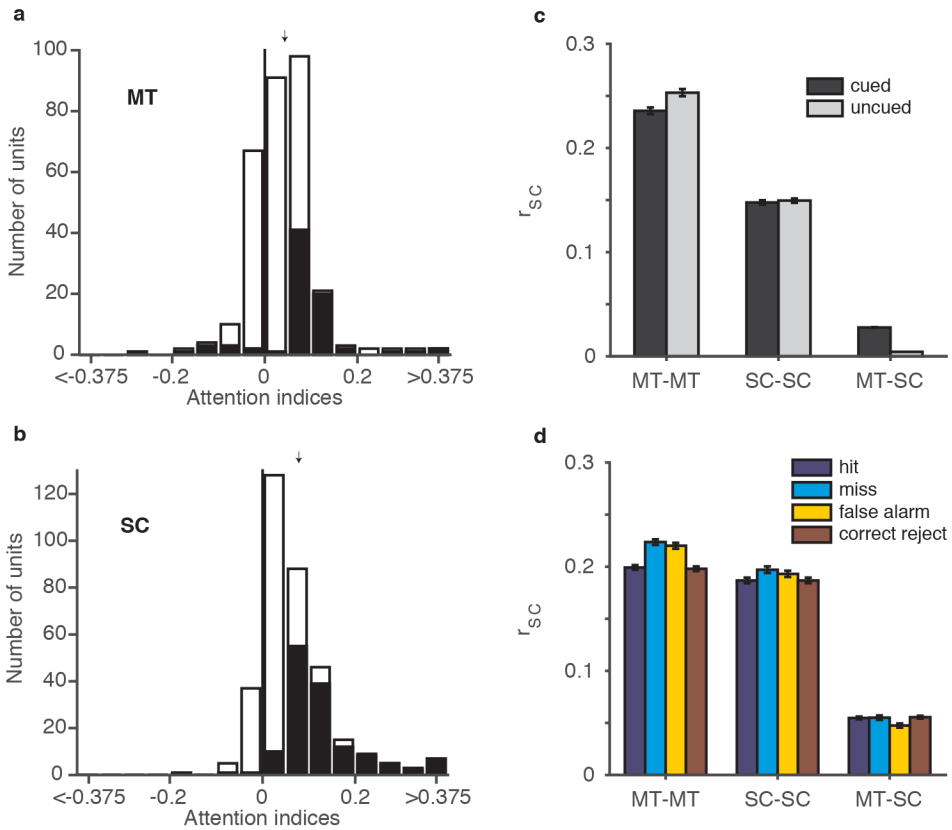

Supplemental Figure 1. Effects of attention on common analyses of individual units and pairs of units (A) Attention increases firing rates in MT, quantified as the difference in firing rates in the different attention conditions divided by the sum. Units with significant differences in average responses for the two conditions are specified by black bars (t-test, p<.05). This distribution (mean = 0.04, median = 0.04) is significantly different from zero (Wilcoxon signed rank test, $p< 10^{-21}$). (B) Same as A, for SC data. This distribution (mean = 0.073, median = 0.05) is significantly different from zero (Wilcoxon signed rank test, $p< 10^{-43}$). (C) Within and between area noise correlations calculated from spike counts during stimulus presentations that preceded successful maintenance of fixation from trials that ended with either a hit or miss or were a successful catch trial. Attention decreases average correlations within MT (Wilcoxon signed rank test, $p< 10^{-12}$), not in the SC (Wilcoxon signed rank test, p=0.8) and increases them between the two areas (Wilcoxon signed rank test, $p< 10^{-41}$). Error bars are standard error of the mean. (D) Within and between area noise correlations calculated from spike counts that immediately preceded different behavioral outcomes during cued trials. Misses and false alarms are associated with higher correlations within MT (t-test, $p<10^{-3}$) and SC (t-test, $p<10^{-3}$) but not between the two areas (t-test, p=0.23). Error bars are standard error of the mean.

Supplemental Figure 2

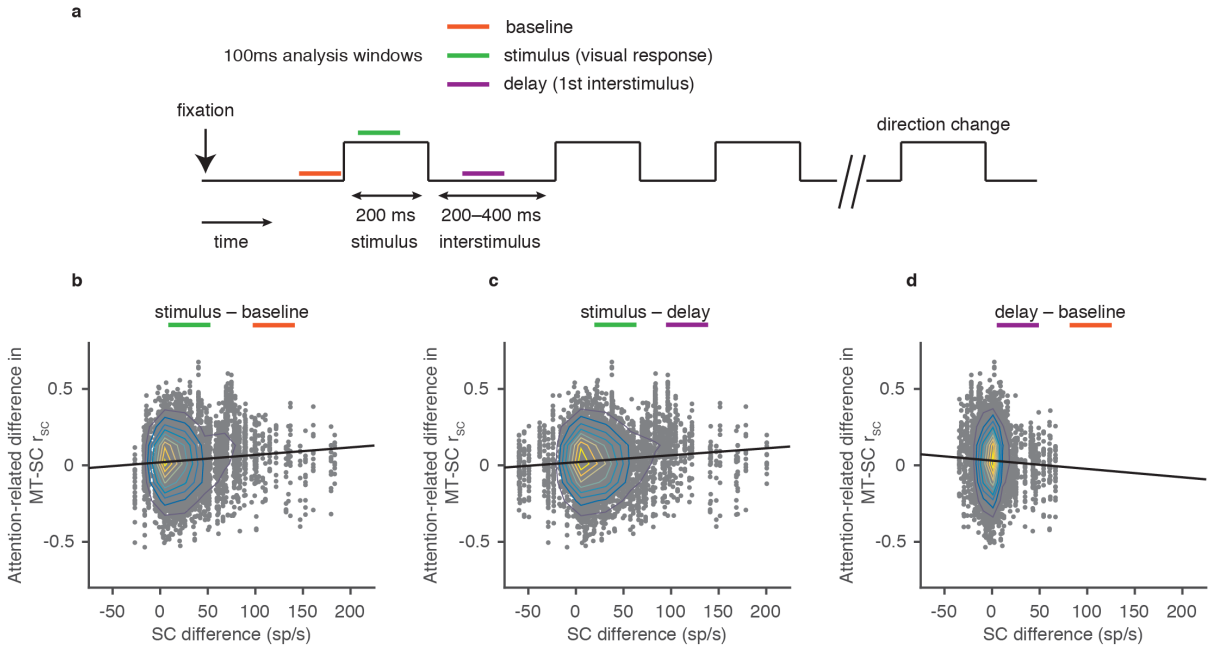

Supplemental Figure 2. Relationship between SC responses during different task epochs and attention-related correlation changes with MT. (A) Schematic of task timing depicts the three 100ms epochs used to count spikes in SC units. The baseline period began 100ms before the first stimulus appeared, which is after stable fixation had been acquired. The stimulus period was shifted 30 ms after the appearance of the visual stimulus, to account for the earliest visual latencies observed in the SC. The delay period began 100ms after the first stimulus turned off and always ended prior to the onset of the second stimulus. (B) Attention-related changes in MT-SC r_{SC} plotted against the difference between each SC unit's response during the stimulus and baseline periods. There are multiple MT-SC correlation differences measured for each SC unit. Correlations between MT and SC were calculated using the same data and methods as Supplementary Figure 1C (Pearson correlation, rho=0.087, $p< 10^{-12}$). Isolines depicting the decile boundaries are overlaid over the individual data points. (C) Similar to B, but data are now sorted by the difference between each SC unit's response during the stimulus and delay periods (Pearson correlation, rho=0.092, $p< 10^{-14}$). (D) Similar to B, but data are now sorted by the difference between each SC unit's response during the delay and baseline periods (Pearson correlation, rho=-0.042, $p< 10^{-4}$).