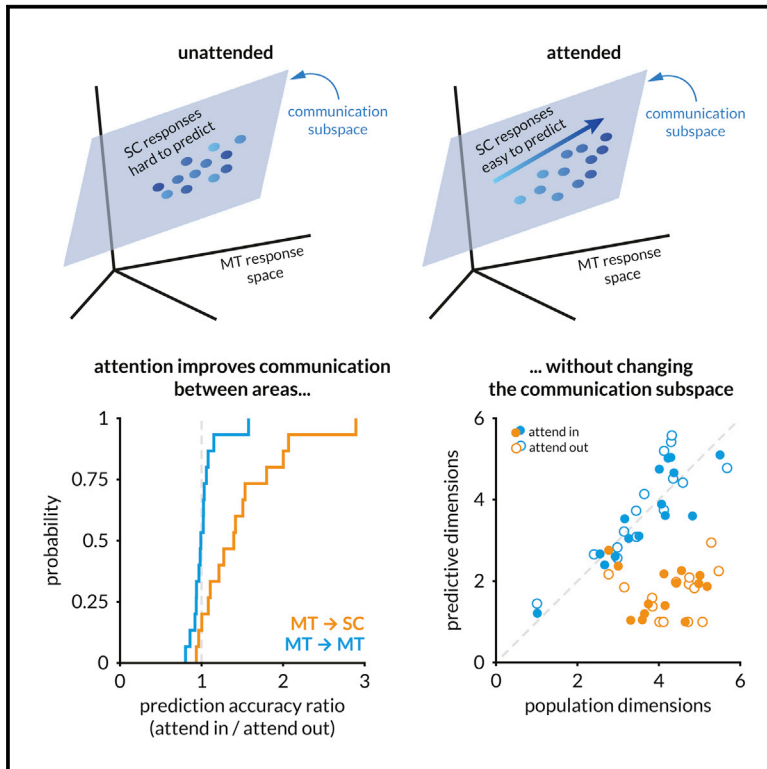


Current Biology

Attention improves information flow between neuronal populations without changing the communication subspace

Graphical abstract



Authors

Ramanujan Srinath, Douglas A. Ruff, Marlene R. Cohen

Correspondence

ramanujan@pitt.edu

In brief

Srinath et al. study how attention affects information flow between two areas. Attention improves prediction accuracy of linear models between areas MT and SC, but not within areas. Therefore, attention could affect decision making by preserving local sensory information processing and altering information shared with other areas.

Highlights

- Attention improves linear response prediction between MT and SC
- Prediction accuracy within areas is not affected
- The dimensionality of the communication subspace between MT and SC remains the same
- Improvements in prediction accuracy are not contingent on dimensionality changes



Article

Attention improves information flow between neuronal populations without changing the communication subspace

Ramanujan Srinath,^{1,2,*} Douglas A. Ruff,¹ and Marlene R. Cohen¹¹Department of Neuroscience and Center for the Neural Basis of Cognition, University of Pittsburgh, Pittsburgh, PA, USA²Lead contact*Correspondence: ramanujan@pitt.edu<https://doi.org/10.1016/j.cub.2021.09.076>

SUMMARY

Visual attention allows observers to change the influence of different parts of a visual scene on their behavior, suggesting that information can be flexibly shared between visual cortex and neurons involved in decision making. We investigated the neural substrate of flexible information routing by analyzing the activity of populations of visual neurons in the medial temporal area (MT) and oculo-motor neurons in the superior colliculus (SC) while rhesus monkeys switched spatial attention. We demonstrated that attention increases the efficacy of visuomotor communication: trial-to-trial variability in SC population activity could be better predicted by the activity of the MT population (and vice versa) when attention was directed toward their joint receptive fields. Surprisingly, this improvement in prediction was not explained by changes in the dimensionality of the shared subspace or in the magnitude of local or shared pairwise noise correlations. These results lay a foundation for future theoretical and experimental studies into how visual attention can flexibly change information flow between sensory and decision neurons.

INTRODUCTION

Perhaps the most impressive hallmark of the nervous system is its flexibility. We effortlessly modulate the extent to which we rely on particular sensory information in different contexts. Visual attention dramatically affects perception and a wide variety of measures of neural activity in essentially every visual and visuomotor brain area (for reviews, see Maunsell¹ and Moore and Zirnsak²). Attention flexibly modulates signatures of neuronal activity, including trial-averaged firing rates,^{1,3,4} shared variability between pairs of neurons in the same^{5–20} and different brain areas,^{15,21–23} interdependence of neuronal populations on a range of timescales,^{15,24–43} and the dimensionality of population activity within each brain area.^{44–46}

The behavioral effects of attention make it clear that visual information can be flexibly routed: a stimulus can either guide or be unrelated to a perceptual decision, depending on the task condition.^{1,47–49} In the visual system, neurons in each area send projections to a variety of different sensory, association, and motor areas, and only a small proportion of neuronal population activity is shared between even highly connected brain areas.⁵⁰ Recent work used correlative methods to identify a functional “communication subspace,” which consists of the dimensions of neuronal population space in which trial-to-trial variability is shared between areas.^{50,51} We similarly adopt the term “communication” to refer to functional communication (i.e., shared trial-to-trial variability in responses to the same visual stimulus).

An exciting possibility is that modulations in the shape or the constitution of this subspace could be a substrate for flexible,

attention-dependent routing of sensory information. Compared to its behavioral effects, attention has remarkably modest effects on the amount of visual information encoded in visual cortex.¹⁴ Instantiating task or attentional flexibility via flexible routing rather than information coding could allow the brain to retain irrelevant visual information for future behavior or memory while the most relevant visual information guides behavior.

We considered three potential, non-mutually exclusive means by which information might be flexibly shared between visual cortex and premotor neurons involved in decision making. Attention might modulate information flow between areas by (1) changing the way visual or task information is represented in a local population, (2) changing the communication subspace itself, and/or (3) changing the efficacy of information transfer (Figure 1D).

Our strategy was to analyze functional communication between neuronal populations in visual and premotor areas while animals switched attention toward or away from their joint receptive fields. We recorded simultaneously from dozens of visual neurons in the medial temporal area (MT) and oculo-motor neurons in the superior colliculus (SC) with overlapping receptive fields while rhesus monkeys performed a task in which they switched spatial attention, alternately biasing decisions toward or away from the stimulus in the joint receptive fields of the recorded neurons. We used recently published methods for analyzing functional relationships between populations of neurons by assessing the dimensionality of shared variability and the extent to which the activity of one population could be predicted by the other.^{50,51} We focused on trial-to-trial fluctuations



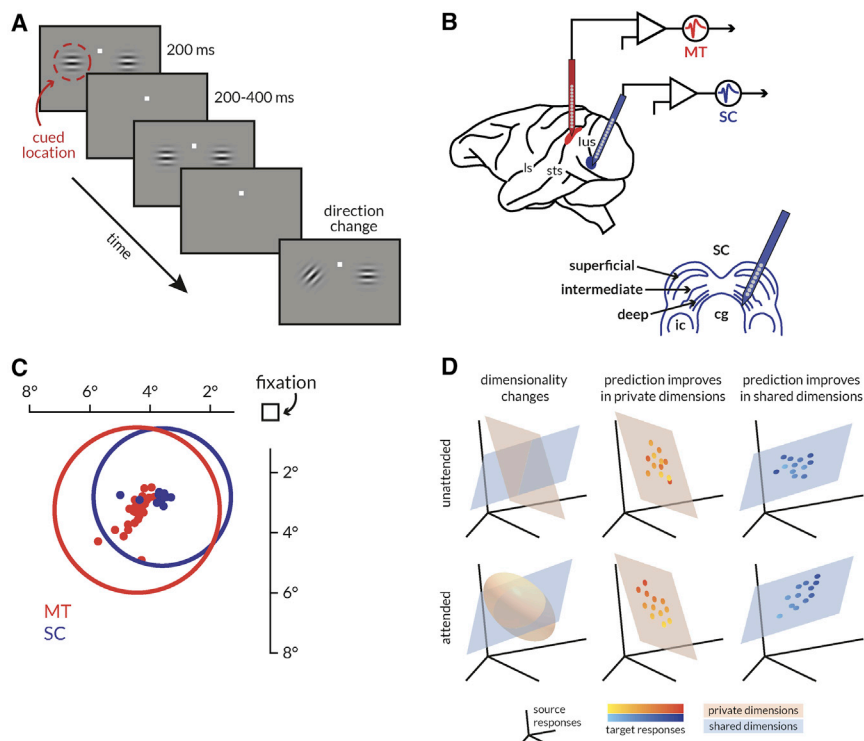


Figure 1. Behavioral task, recording sites, receptive fields, and schematic of hypotheses

(A) Schematic of the motion direction change detection task. The monkeys were cued in blocks of trials to expect changes in motion direction at one of two spatial locations (cue was 80% valid). The monkey started the trial by fixating a central spot. Two small Gabor stimuli synchronously flashed on for 200 ms and off for a randomized period of 200–400 ms. One of the stimuli was positioned inside the joint receptive fields of the MT and SC neurons, and the other was placed in the opposite hemifield. Both stimuli moved in a direction that was chosen to drive the MT population well. After a randomized number of stimulus presentations (between 2 and 13), the direction of one of the stimuli changed. The monkeys were rewarded for making a saccade to the direction change in either location. We analyzed neuronal responses to all identical stimulus presentations except the first to minimize the effect of adaptation. (B) Illustration of recording locations. Populations of MT and SC neurons were recorded with linear 24-channel moveable probes from the right hemisphere of two monkeys as they were doing the behavioral task described in (A).

(C) Receptive field locations of recorded units from an example recording session. The dots represent the receptive field centers of 28 MT (red) and 26 SC (blue) units. The circles represent the size and location of the median receptive field from each area.

(D) Schematics describing the hypotheses about attention-related changes in information flow between two areas. Each icon depicts the response space of the source area (the responses of the first n neurons or principal components, for instance) and orange and blue surfaces that represent two subspaces for the private or shared fluctuations in neural activity, respectively. The two rows of icons represent the attended and unattended conditions (when attention was directed toward or away from the receptive fields of the recorded neurons), and each column describes the expected result of each of the following hypotheses. Left: attention could alter the dimensionality of the private, shared, or both subspaces. If attention only modified local representations, then the number of private dimensions that explain the local neural fluctuations would change. Middle: alternatively, attention could modulate information flow by enhancing or diminishing the extent to which neural activity in a target population tracks the neural activity of its source. If attention acted via this mechanism locally, then prediction would improve in private dimensions. Right: if attention modulated functional communication by modulating information flow across areas, then prediction would improve in shared dimensions.

in responses to the same visual stimulus because these are related to functional connectivity rather than simply reflecting tuning for similar stimuli (for review, see Cohen and Kohn⁵² and Umakantha et al.⁵³). Furthermore, fluctuations on the timescale of single stimulus presentations^{13,54} or longer⁴⁴ have been shown to be correlated with choice behavior.

The three hypothesized ways that attention could change information routing (Figure 1D) make different predictions about how attention changes the relationship between trial-to-trial fluctuations in MT and the SC. The first hypothesis, that attention changes the local representation of visual or premotor information, suggests that attention should change the dimensionality and/or information content of the representation within an area. The second hypothesis, that attention changes the communication subspace, predicts that attention changes the dimensionality of the subspace and/or the identity of the dimensions it includes. Either of these changes could affect information flow: increasing or decreasing the number of shared dimensions could widen or restrict shared information, while changing the identity of the shared dimensions could change the nature of the shared information. The third hypothesis, that attention improves the efficacy of functional communication between visual

and premotor neurons, predicts that attention enhances the extent to which we can predict the responses of neurons in one area from neuronal responses in the other area.

We found strong evidence for our third hypothesis. Trial-to-trial variability of the population of SC neurons was better predicted by the activity of MT neurons (and vice versa) when attention was directed inside their joint receptive fields. This enhanced functional communication between areas could not be explained by increases or decreases in the raw magnitude of private or shared pairwise noise variability or a change in the number of private or shared dimensions of neuronal population activity.

This enhanced functional communication was not restricted to interactions between visual and premotor neurons. The effects of attention on functional communication were similar between MT and visual or motor neurons in the SC and in an independent dataset of simultaneously recorded neurons in primary visual cortex (V1) and in MT. Even though the attention-related change in pairwise correlations and response dimensionality within V1 was small compared to MT or SC, attention significantly enhanced our ability to predict the responses of single MT neurons from population activity in V1.

Our study provides a blueprint for combining behavioral paradigms that vary cognitive processes with dimensionality reduction and regression analyses to study how information can be flexibly routed in the nervous system. These results are the first demonstration of how attention affects the activity of distinct but connected populations of neurons in a way that affects the functional interactions between visual and premotor neurons. They suggest a set of possible mechanisms (see [Discussion](#)) by which cognitive processes can affect perceptual decision making in ways that are independent of changes to the local neuronal representations.

RESULTS

We compared evidence consistent with several potential mechanisms for flexible routing of information. We chose a widely studied cued direction change detection task to study the behavioral effects of attention on visual perception and three brain regions that are known to contribute to motion perception and visually guided decision making—primary visual cortex (V1), the MT, and the SC. Rhesus monkeys performed the motion change detection task described in [Figure 1A](#). Briefly, after the monkey fixated a central spot, two peripheral, full-contrast, drifting Gabor stimuli synchronously flashed on (for 200 ms) and off (for a randomized period between 200 and 400 ms) until, at a random, unsignaled time, the direction of one of the stimuli changed from that of the preceding stimuli. The monkey received a liquid reward for directing a saccade to the stimulus that changed within 450 ms of its onset. Using instruction trials prior to each block, attention was cued in blocks between each of the visual stimuli. While rhesus monkeys performed the task ([Figure 1A](#)), we recorded simultaneously from either dozens of neurons in MT and SC ([Figure 1B](#)) with overlapping receptive fields (RFs) ([Figure 1C](#); different aspects of these data were previously reported in Ruff and Cohen¹⁴) or from several dozen neurons in V1 and a single MT neuron ([Figure 7](#); different aspects of these data were previously reported in Ruff and Cohen^{15,55}). This stimulus was placed either inside the overlapping receptive fields of the recorded neurons or in the opposite hemifield ([Figure 1C](#)). Throughout this manuscript, “attend in” refers to the trials where attention was directed toward the joint RFs and “attend out” refers to trials where attention was directed to the opposite hemifield. The monkey was rewarded for making a saccade to the location of the direction change, which occurred at a random and unsignaled time. The ability of the animal to detect the median difficulty changes in grating direction is enhanced by ~25% on average across sessions when attention was directed to the location of the change (cued 76.5% detected; uncued 51.8% detected).¹⁴ We analyzed the spike counts of each visually responsive multi-unit recorded from MT and SC during presentations of identical Gabor stimuli before the direction change (excluding the first presentation in each trial to remove adaptation effects). We also analyzed spike counts of each SC unit with elevated firing rates before saccade onset to the contralateral visual field. In the V1-MT dataset, we tested our hypotheses on the responses of groups of V1 neurons whose receptive fields overlapped either of two small stimuli, both of which were inside the RF of the MT neuron.^{15,55}

Signatures of population interactions that underlie attentional mechanisms

We tested the following non-mutually exclusive hypotheses (schematized in [Figure 1D](#)) about how attention might modulate information flow within and between areas. (1) Attention primarily modulates communication between areas by changing the dimensionality of either the private or the shared subspace ([Figure 1D](#), left column). (2) Attention improves the fidelity of communication within local populations; this would be observable as an improvement in the ability to predict the activity of one subset of neurons in a population from the activity of a different subset of neurons in the same area ([Figure 1D](#), middle column). (3) Attention improves the fidelity of communication across brain regions; this would be evident in the improved accuracy of prediction of neural activity of one region using the activity of the other and vice versa ([Figure 1D](#), right column).

Prediction of SC activity from MT activity using linear models improves with attention

Testing the predictions of our hypotheses requires calculating the ability to predict the activity of one population of neurons from another and identifying the dimensions of neural population space through which functional communication occurs. We plot the results of these analyses for one representative session in [Figure 2](#). We used ridge regression to impose a sparse mapping between random subsets of MT neurons and the full populations of SC neurons in each attention condition (see [STAR Methods](#) and Semedo et al.⁵⁰).

Several features of this example recording session were typical of our dataset. First, no subset of the recorded MT neurons could predict SC neural activity as well as the full population; the prediction accuracy monotonically increased with the addition of MT neurons. Second, the accuracy of prediction was significantly improved in the attend in trials versus attend out trials across all sub-selections of the MT population. Third, attention also improved the ability to predict random subsets of SC neurons from the full population of recorded MT neurons ([Figure 2B](#)).

To determine the relationship between these measures of functional communication between neuronal populations in MT and the SC and more well-studied effects of attention, we next calculated traditional metrics neuronal activity like the average pairwise spike count correlation ([Figure 2C](#)) and population firing rate ([Figure 2D](#)). For this session, attention significantly decreased spike count correlations in both MT and SC but did not have an effect on variability shared between pairs of neurons in different brain areas. Attention also significantly increased mean firing rates in this session. Firing rate and correlation changes across sessions are detailed in [Figures S1](#) and [S3](#).

For the example session, we observed no attention-related change in the population dimensionality in MT (~5 dimensions; [Figure 2E](#)) and SC (~3.5 dimensions; [Figure 2G](#)), defined as the smallest number of dimensions that captured 95% of the variance in the shared covariance matrix (assessed using factor analysis;⁵⁶ also see [STAR Methods](#) for code and other resources).

We next tested whether, as between two areas of visual cortex,⁵⁰ interactions between MT and the SC are limited to a subset of dimensions of neural population space. For the example

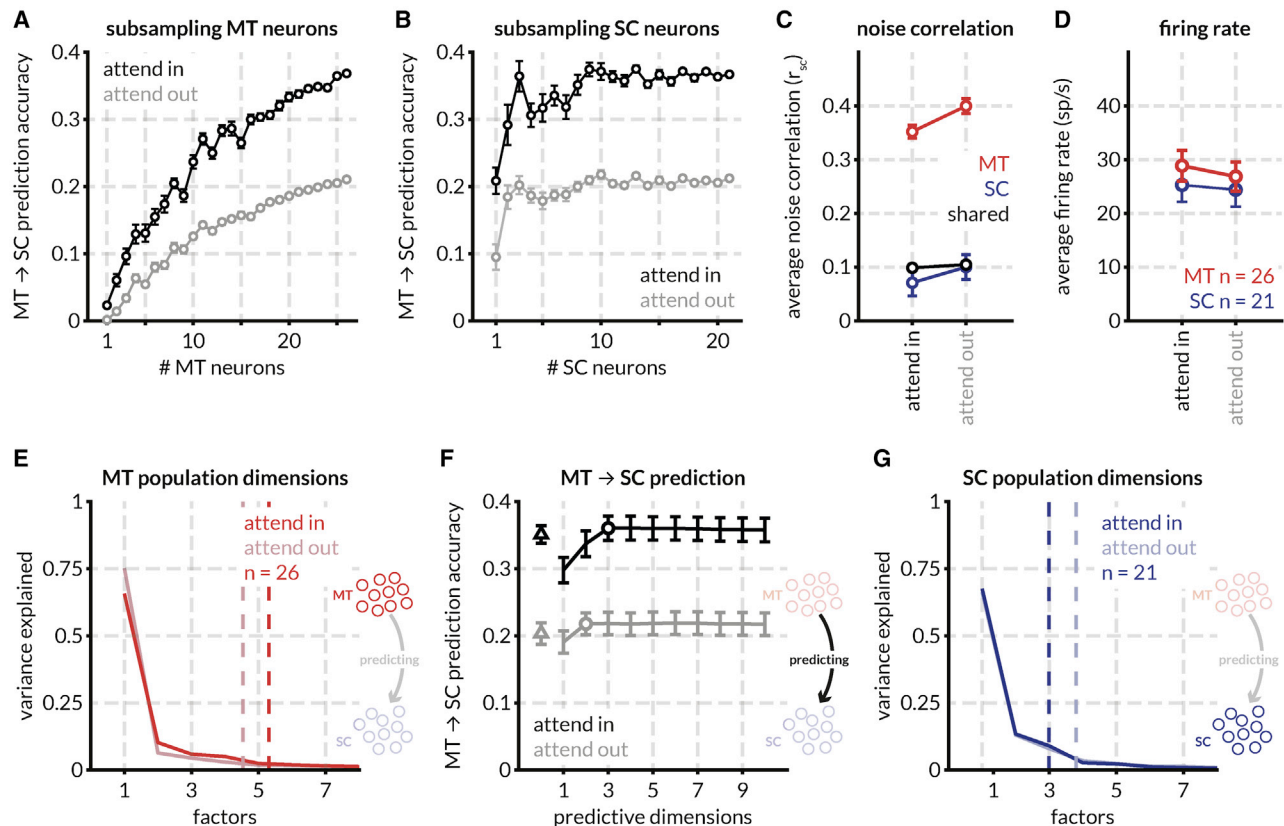


Figure 2. Attention improves prediction of SC activity from MT activity, increases firing rate, and decreases spike-count correlations in an example recording session

(A) For an example session, the prediction accuracy of 1–26 randomly sampled (without replacement) MT neurons predicting the activity of a population of 21 SC neurons in the two attention conditions (attend in refers to the trials in which attention was directed within the joint RFs of the MT and SC neurons and attend out refers to trials in which attention was directed in the opposite hemifield). Prediction was performed using a linear model with ridge regression, and prediction performance was defined as 1-cvLoss, where cvLoss is the average cross-validated normalized square error (NSE) for the smallest ridge parameter for which the performance was within 1 SEM of the peak performance. Each point represents the mean prediction performance for n MT neurons predicting the full population of SC neurons. Error bars represent the SEM across random subsamples of n neurons.

(B) Same as (A) but for predicting random subsets of SC neurons using the activity of the full population of MT neurons, showing that the effect of attention on MT-SC communication is not limited to a subpopulation of either the MT or SC neurons recorded in this session.

(C) Spike count correlation (r_{sc}) defined as the correlation between the responses of pairs of neurons to all stimulus presentations for all MT neurons (325 pairs, red), SC neurons (210 pairs, blue), and MT-SC pairs (546 pairs, black). Attention decreases spike count correlations in MT ($p = 1.2 \times 10^{-10}$; Wilcoxon signed-rank test [WSRT]) and SC ($p = 0.0206$; WSRT) but has no effect on pairwise correlations across areas ($p = 0.2$; WSRT) for this recording session. Note that the decrease in average SC neuron pairwise correlations in this session is not representative of the trend across sessions. See Figure S1 for r_{sc} for all pairs across recording sessions.

(D) Neuronal firing rates increase with attention in MT ($p = 8.3 \times 10^{-6}$; WSRT) and SC ($p = 0.04$; WSRT) for this session. See Figure S1 for firing rates for all neurons across sessions.

(E) Factor analysis of MT population responses for this session reveals that 90% of the variance in the MT response fluctuations can be accounted for by ~5 dimensions. The number of population dimensions is greater for the attend in condition versus the attend out condition. The arrow in the icon signifies that the MT population (source) is being used to predict the SC population (target), henceforth labeled as MT \rightarrow SC prediction.

(F) Predicting SC activity from MT responses using reduced-rank regression (RR regression) (black and gray lines) and ridge regression (triangle) reveals that the prediction performance for a matched number of trials is dramatically better for the attend in condition (black) versus the attend out condition (gray). The optimum number of dimensions (circle) for the RR regression was defined as the lowest number of dimensions for which prediction performance was within 1 SEM of peak performance. This performance is at least as good as the performance of the ridge regression performance that uses all the source dimensions for prediction (the difference between the RR regression prediction and the ridge regression prediction was not significant across sessions; data not shown). The number of source dimensions required for optimum regression performance was 3 for attend in and 2 for attend out, suggesting that fewer dimensions are required for communication between MT and SC than the total number of population dimensions.

(G) Factor analysis of SC neurons reveals that 90% of the variance in the SC response fluctuations can be accounted for by 3 to 4 dimensions. For this session, the number of population dimensions is greater for the attend out condition versus the attend in condition.

session in Figure 2, only 2 to 3 dimensions of MT activity (identified using reduced rank regression; STAR Methods; defined at the number of dimensions at which the curves in Figure 2F reach

asymptote) predicted SC activity at least as well as a full linear model (fit using ridge regression; STAR Methods). The prediction accuracy for the attend in trials was significantly better than the

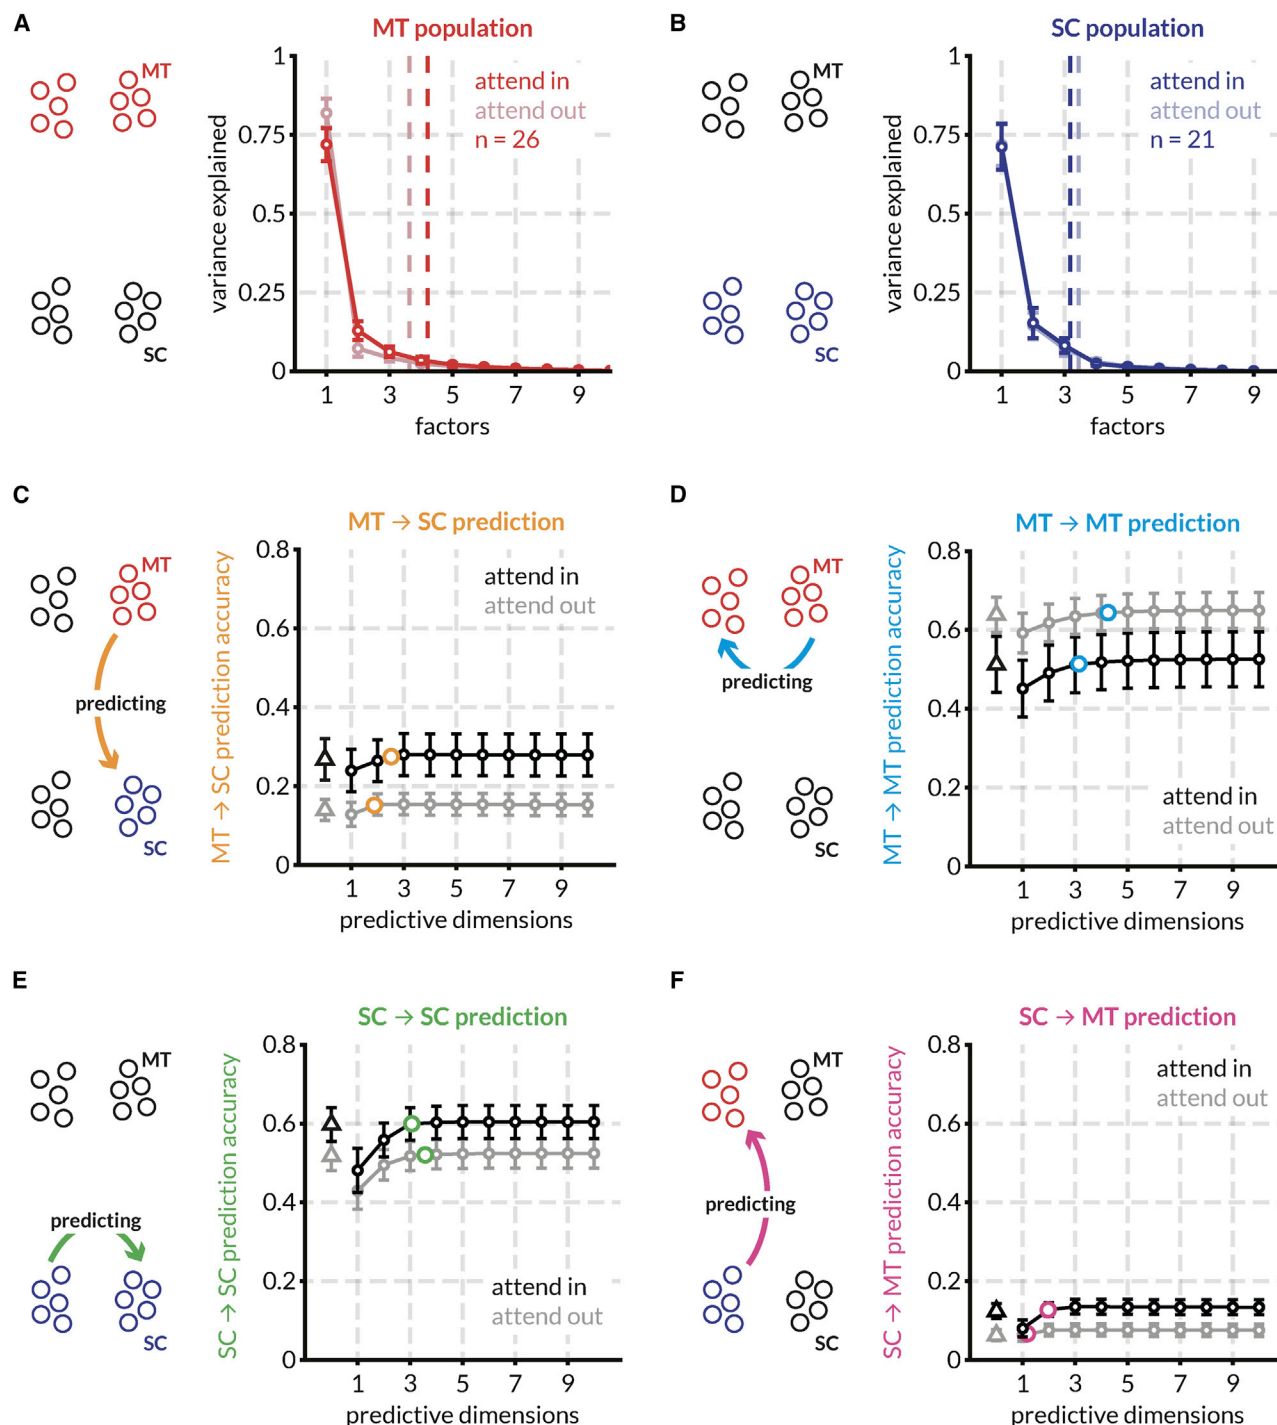


Figure 3. Randomly partitioned populations of MT and SC neurons predict activity within and across areas better with attention for the same example session

To compare prediction performance for inter- and intra-areal interactions, we randomly split both the populations of MT and SC neurons into two halves each—the target and source halves—as indicated in the icons. Each source half was used to predict the activity of both target halves using both the full linear model (ridge regression) and the RR regression model. This split was done 20 times, and the mean performance across the random splits is shown in (C)–(F). Error bars indicate the SEM across these splits.

(A) Factor analysis of MT neurons reveals that 95% of the variance in the MT response fluctuations can be accounted for by ~4 dimensions on average across all splits for this session. The number of population dimensions is greater for the attend in condition versus the attend out condition.

(B) Same as (A) for SC neurons. For this session, SC population fluctuations are captured by ~3 dimensions in both attention conditions.

(legend continued on next page)

attend out trials irrespective of the number of predictive dimensions (the black line is always above the gray line in Figure 2F).

Attention improves prediction accuracy for inter-areal communication channels

Testing our hypothesized mechanisms of information flow (Figure 1) requires determining how attention affects the dimensionality and informativeness of interactions within and between populations of neurons in MT and the SC. We therefore fit linear models for repeated random splits of the populations of recorded MT and SC neurons in all four directions—MT \rightarrow SC, MT \rightarrow MT, SC \rightarrow MT, and SC \rightarrow SC. We depict the effect of attention on these four communication channels (for the same single session as in Figure 2) in the form of mean prediction accuracy across all tested population splits (Figures 3C–3F). For this session, the prediction performance improves with attention for all functional communication channels except within MT, where it depreciates. We estimated the population dimensionality of each of the randomly split populations of MT and SC neurons using factor analysis to compare with number of predictive dimensions (Figures 3A and 3B). Consistent with results for the full population above, the number of dimensions within each area is not affected by attention.

Across sessions, prediction performance between MT and the SC improves with attention without changing the dimensionality of that communication (Figure 4). Whereas prediction accuracy for intra-areal communication was consistently high and remained unaffected by attention, the prediction accuracy for inter-areal communication significantly improved with attention (Figure 4B, which shows the ratios of the number of predictive dimensions and of the prediction accuracy in the two attention conditions). Attention does not affect the number of predictive dimensions required for communication within and across areas (the marginal distributions of ratios are centered at and not significantly different from 1) but improves the prediction accuracy between MT and the SC (the distributions of ratios of MT \rightarrow MT and SC \rightarrow SC prediction accuracy are centered at and not significantly different from 1, but the ratios of MT \rightarrow SC and SC \rightarrow MT prediction accuracy are significantly greater than 1; see also the distributions for each communication channel in Figures S2 and S3). We analyzed the effect of attention on the population dimensionality and the number of dimensions required for prediction in detail in the next section. To quantify the size of the effect of attention for each session, we calculated a modulation index for prediction performance (difference of prediction performance in the two attention conditions divided by the sum). The inter-areal modulation indices (MT \rightarrow SC: 0.164 ± 0.04 and SC \rightarrow MT: 0.124 ± 0.04) were considerably larger than intra-

areal modulation indices (MT \rightarrow MT: 0.005 ± 0.019 and SC \rightarrow SC: 0.068 ± 0.02 ; Figure S2). Importantly, while the prediction performance of SC \rightarrow SC regressions across sessions was significantly greater for the attend in than the attend out condition, the effect size was small. A larger dataset of neurons and sessions may be required to definitively test the hypothesis that attention changes the efficacy of within area communication in oculo-motor areas.

While attention is known to affect the mean pairwise spike count correlations within and between areas,^{6,11,15,16} we found that attention-related improvements in inter-areal prediction accuracy are not contingent on increases or decreases in mean spike count correlations between pairs of neurons within MT, within SC, or between MT and SC across sessions (Figure S3). We found that the mean spike count correlations between pairs of MT neurons and pairs of SC neurons are related to intra-areal prediction accuracy within MT and SC populations, respectively, but this relationship does not extend to mean spike count correlations between MT-SC pairs and inter-areal response predictions.

We further investigated the relationship between spike count correlations and prediction performance by examining correlations that may persist on timescales longer than a single stimulus presentation. We time shifted responses of the target population by 0–5 stimulus presentations and recalculated prediction accuracy and mean spike count correlations as described above. As expected, both spike count correlations and prediction accuracy within and across brain areas decay with time shift (shown for a single session in Figures 5A–5D). To visualize this decay across sessions and to directly compare spike count correlations and prediction accuracy, we normalized each decay profile by the value at time shift 0 and plotted the median decay profile across sessions (Figures 5E and 5F). Across sessions, whereas all pairwise correlations decay to ~ 0 within 3 to 4 time-shifted stimulus presentations, prediction performance for MT \rightarrow SC, SC \rightarrow MT, and SC \rightarrow SC predictions does not. Parallel to the results in Figure 4, absolute MT \rightarrow SC prediction performance is higher for attend in trials compared to attend out trials for all time shifts, and performance for attend in trials decays to $\sim 50\%$ of peak value within 5 time shifts and to $\sim 75\%$ of peak value for attend out trials, i.e., performance remains relatively high across long timescales. This analysis emphasizes the differences between mean spike count correlations and prediction accuracy, and it also points to potential neuromodulatory mechanisms that may underlie the attentional effects on the same timescales.

The connectivity and functional roles of populations of SC neurons differ by layer, so we made use of our recordings that spanned layers to investigate whether functional communication

(C) Average prediction performance for the full model (black and gray triangles) and the RR regression model (black and gray circles) across random splits of the MT and SC populations. The orange circle indicates the average optimum performance and average number of optimum prediction dimensions across the random splits. For each session, this point of optimum performance is plotted in different comparisons in the following figures. For this session, attention improves MT \rightarrow SC prediction performance. For all predictions, the RR regression model performs at least on par with the full linear model using ridge regression.

(D) Same as (C) for MT \rightarrow MT predictions. For this session, attention degrades prediction performance. The average optimum performance and average optimum prediction dimensions are indicated with blue circles.

(E) Same as (C) for SC \rightarrow SC predictions. For this session, attention improves prediction performance. The average optimum performance and average optimum prediction dimensions are indicated with green circles.

(F) Same as (C) for SC \rightarrow MT predictions. For this session, attention improves prediction performance. The average optimum performance and average optimum prediction dimensions are indicated with pink circles.

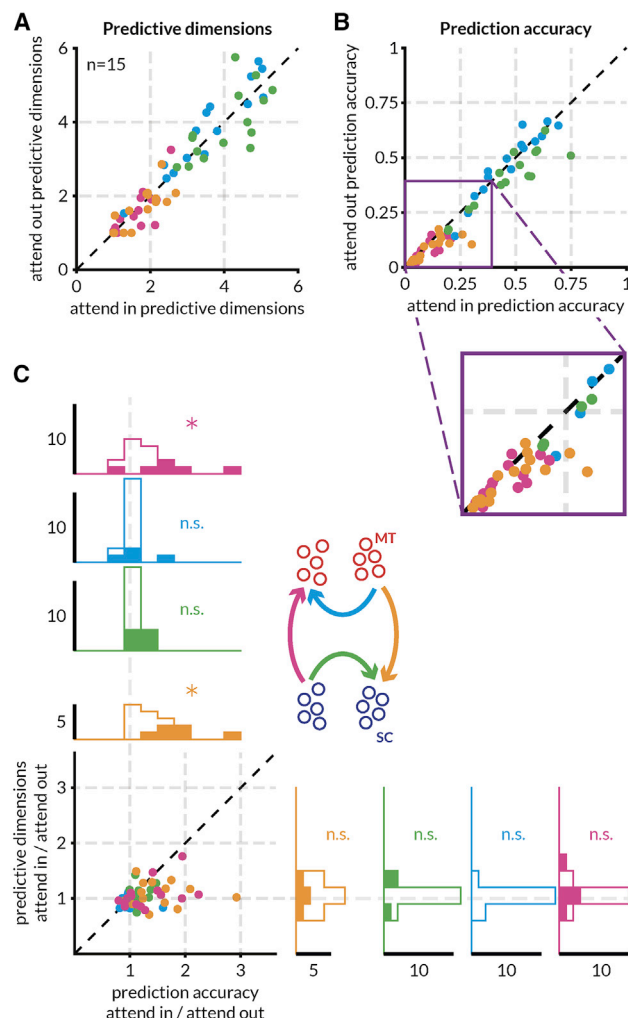


Figure 4. Attention improves the accuracy of across-area prediction, but not within-area prediction, without altering the dimensionality of the communication subspace

Each point of a given color represents a recording session. The color scheme is depicted in the icon in (C) and is consistent with other figures.

(A) Attention does not affect the dimensionality of the interaction between MT and SC neurons. Each point represents the average number of optimum predictive dimensions for each session for one of the four predictions—MT → SC (orange), MT → MT (blue), SC → SC (green), and SC → MT (pink)—for the two attention conditions. There was no significant difference between the number of predictive dimensions for any of the four predictions. See Figure S5 for a detailed version of this panel (MT-MT mean 3.67, range 1.5–5.2 for attend in and mean 3.74, range 1.1–5.3 for attend out; SC-SC mean 4, range 2.9–5.3 for attend in and mean 3.9, range 2.85–5.7 for attend out; MT-SC mean 1.8, range 1–2.5 for attend in and mean 1.75, range 1–2.7 for attend out; SC-MT mean 1.6, range 1–2.7 for attend in and mean 1.55, range 1–3.15 for attend out).

(B) Attention significantly increases the prediction accuracy of inter-area, but not intra-area, interactions. Each point represents the average prediction performance across random splits for one of the four predictions. The purple inset affords a zoomed-in view of the relevant part of the plot, which reveals that the points corresponding to the MT → SC (orange) and SC → MT (pink) predictions lie below the unity line. See Figure S2 for a detailed version of this panel. The drop in the absolute values of prediction accuracies compared to Figures 2 and 3 is due to the split source and target populations to enable comparisons of within- and across-area prediction.

between MT and the SC depends on layer as well. MT projections to SC predominantly end in the superficial layers in SC^{57,58} (but also see Lock et al.⁵⁹). Tecto-pulvinar projections from the superficial and intermediate layers of SC end in the inferior pulvinar, which in turn projects to extra-striate areas.^{60,61} Also, there is some evidence that extra-striate projecting lateral geniculate nucleus (LGN) neurons do not receive direct retinal input and are dependent on SC projections across all layers for relaying visual information to MT.^{62,63} Given these laminar differences in cortical and thalamic inputs to and outputs from SC, we tested whether there is a difference between the attentional effect on information flow across functional classes of SC neurons. To classify SC neurons, we calculated an oculo-motor score based on SC neuron responses to visual stimuli and responses just prior to saccade onset (STAR Methods), divided each population into two groups based on the rank ordering of oculo-motor scores, and repeated our analyses (Figure S4). We found no significant differences in the effect of attention on either the prediction accuracy or the number of dimensions required for prediction between the SC populations split by oculo-motor score (labeled *visual* and *motor* for brevity). Compared to random splits of the SC population, when split by oculo-motor score, the effect of attention on the prediction accuracy of the SC → SC regression is pronounced (Figure S4C).

Attention does not improve information flow by altering private or communication subspaces

The attention-related improvement in information flow could in principle arise by changing the subspaces of activity responsible for functional communication within or between areas. We did not find evidence that attention changes the dimensionality of any of these subspaces: there was no attention-related change in the dimensionality of the local populations of MT and SC neurons (Figures S5A and S5B, respectively) or in the number of predictive dimensions for the various communication subspaces within and between the two areas (Figures S5C–S5F). We consistently found that more dimensions were required to account for intra-areal communication than to account for inter-areal communication (mean 3.6 for MT → MT and 4 for SC → SC versus 1.8 for MT → SC and 1.6 for SC → MT). This disparity suggests that MT and SC interact via a limited communication subspace.

Semedo et al.⁶⁰ compared the number of predictive dimensions with the number of target dimensions to infer that V1 and V2 communicate using a communication subspace. We first replicated those results in the context of MT-SC communication—when we compared the number of dimensions used for private communication with the number of shared dimensions for MT → MT prediction and MT → SC prediction, we found

(C) The data in (A) and (B) visualized as a ratio of attend in and attend out. The marginal distributions of the ratios of prediction accuracy and predictive dimensions for all four predictions are also displayed. The ratios of means of prediction accuracy for MT → SC (orange) and SC → MT (pink) were significantly greater than 1 ($p = 0.0016$ and $p = 0.012$, respectively; t test). Filled-in histograms indicate sessions for which the mean of the corresponding metric (prediction accuracy or prediction dimensions) in the attend in condition was more than two standard deviations away from the mean of the metric in the attend out condition. The colored arrows in the icon indicate the source and target populations for each of the four predictions.

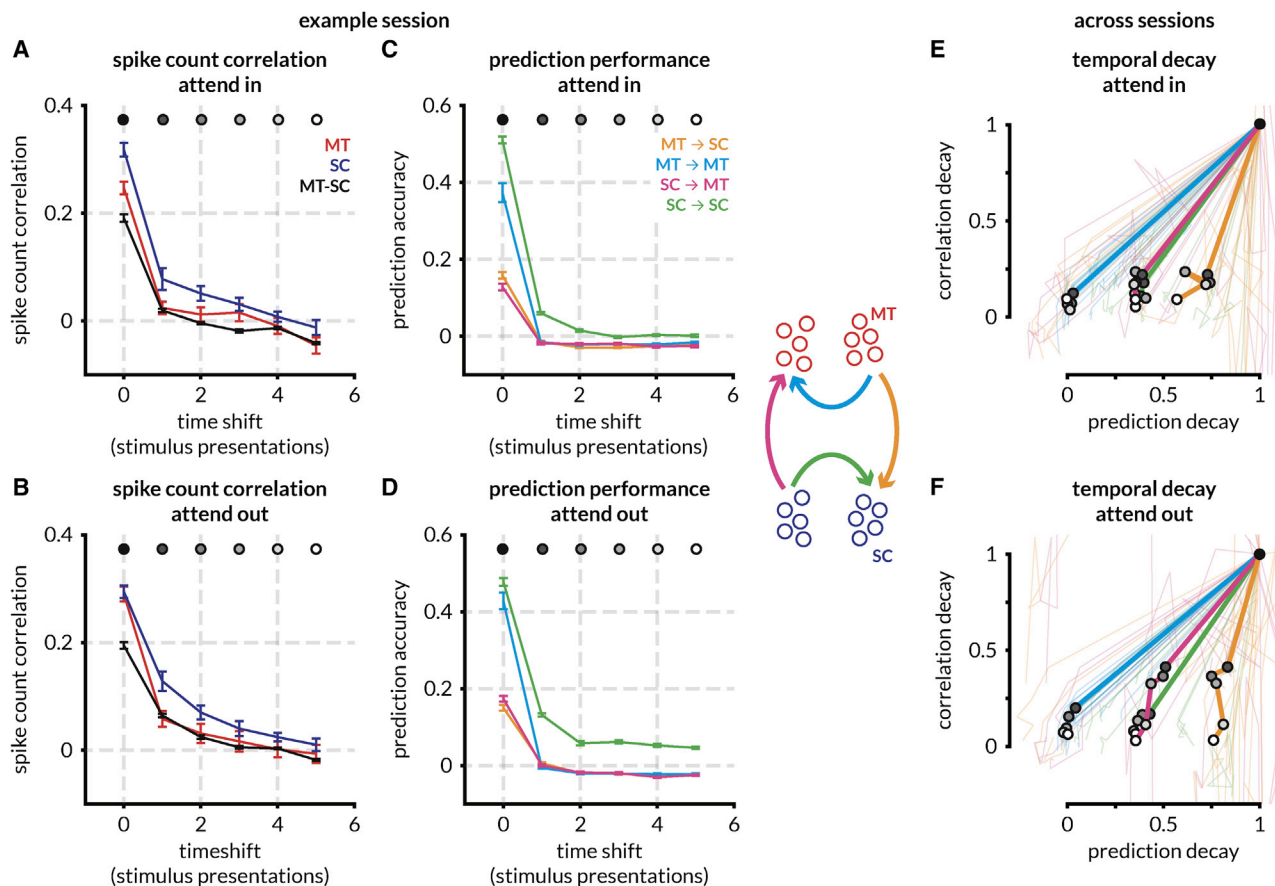


Figure 5. Inter-stimulus (long timescale) spike count correlations within and across areas decay faster than linear predictions

(A)–(D) depict the decay in spike count correlations (y axes in A and B) and in prediction performance (y axes in C and D) within and across areas as a function of time (in stimulus presentations, x axes) for an example session. (E) and (F) depict decay as a proportion of peak values across sessions. Error bars indicate the SEM across MT and SC population splits (see also Figure S3 for comparisons of the effect of attention on prediction accuracy with the effect on spike count correlations and attentional modulation).

(A and B) Spike count correlations decay rapidly across time shifts for MT, SC, and MT-SC neuron pairs for attend in trials. Within- and across-area spike count correlations decay comparatively slower for attend out trials for this example session.

(C and D) Prediction performance rapidly decays to 0 for all predictions in this example session, although within SC prediction performance remains higher than baseline (0) for attend out trials. Note that the profile of decay in the prediction accuracies in this session is not representative of the median across sessions (E and F).

(E and F) Decay in prediction performance (x axis) or pairwise correlations (y axis), normalized by the value at time shift 0 (time-aligned trials between source and target populations). Within-area predictions (MT → MT and SC → SC) are compared with within-area pairwise correlations (MT and SC correlations, respectively), and across-area predictions (MT → SC and SC → MT) are compared with across-area pairwise correlations (MT-SC correlations). Thin lines are decay ratios for each session; thick lines are median values across session for each time shift. Time shifts are indicated with grayscale markers in the median curve.

that significantly fewer dimensions are required for MT → SC communication than are available, but MT → MT communication utilizes all available dimensions (Figure 6A). In addition, we found no difference in the dimensionality of the communication subspace compared to target population dimensionality across attention conditions (Figures 6B and 6C). We found similar results in the SC → MT direction when compared with SC → SC communication (Figures 6D–6F). We found no relationship between the functional communication channels when assessed on a session-by-session basis (Figure S6). We also did cross-prediction analyses (using the attend in linear model to predict attend out data and vice versa) to check whether the structure of the communication subspace changes while keeping its

dimensionality, in turn causing the prediction accuracy to be better (Figure S7). We found that while the intra-areal models performed almost as well when swapped, the inter-areal models suffered a loss in prediction accuracy. This does not necessarily imply that the geometry of the communication subspaces changes with attention but that linear methods are unable to find a common subspace between the two attention conditions (also see Discussion).

Attention improves information flow between V1 and MT

Both MT and SC exhibit relatively large attention-related changes in a number of measures of neuronal activity.^{64–69} Attention-related improvements in information flow may in

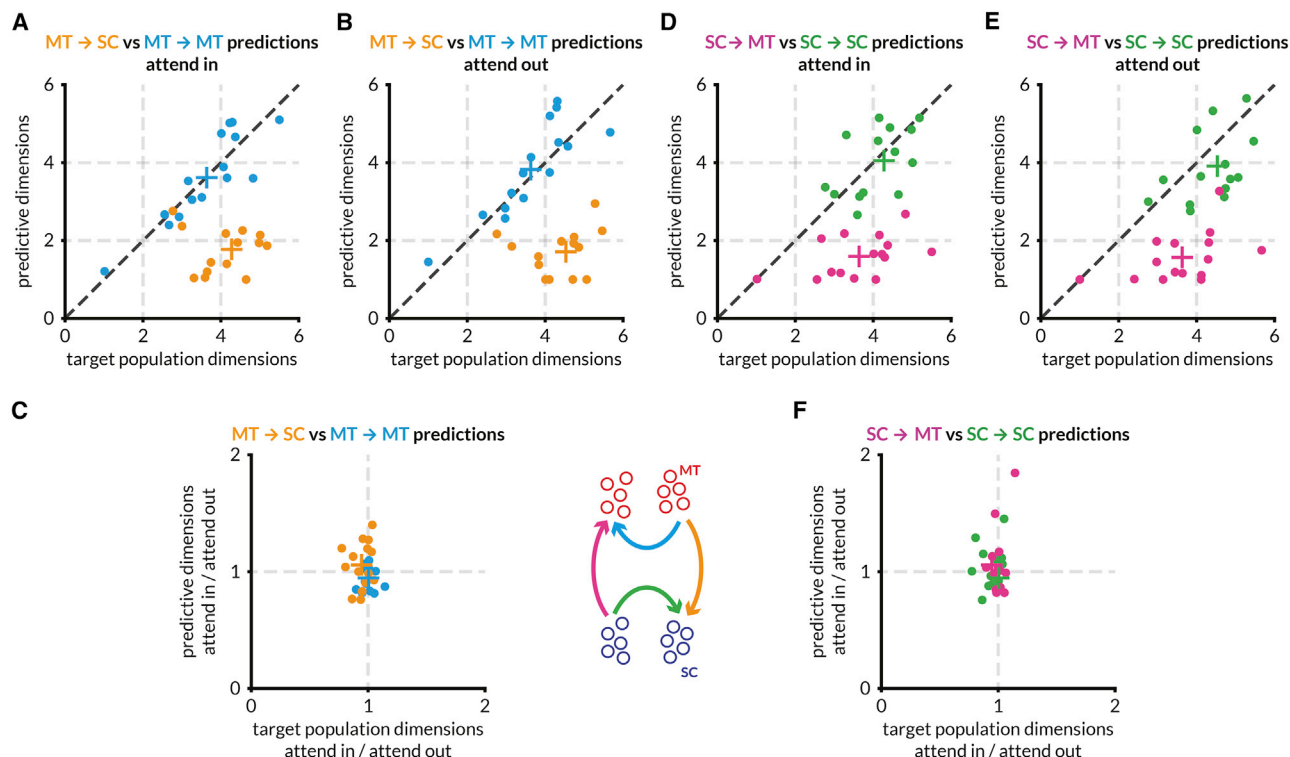


Figure 6. MT and SC populations interact via a communication subspace, but attention has no effect on the dimensionality of the communication subspace

Each point represents a recording session, and the color scheme is the same as other figures. Colored + represents the mean of the corresponding points. This figure compares the number of factors that explain 95% of the variance in the target area (from factor analysis) with the number of dimensions in the source area that are sufficient to predict the target area activity (from RR regression). Qualitative comparisons between the absolute values of the “number of dimensions” from these two analyses are depicted in (A), (B), (D), and (E). The effect of attention is depicted in (C) and (F) (see also Figures S5 and S6 for detailed comparisons of the effect of attention on population and subspace dimensions).

(A) For the attend in condition, the number of private predictive dimensions are greater than the number of shared predictive dimensions in MT. Further, for the MT → SC prediction (orange points), fewer dimensions are required to predict SC activity than are required to explain 95% of the variance in the SC activity, forming a communication subspace in MT that comprises ~2 shared dimensions that are sufficient to predict the ~4-dimensional activity in SC. For the MT → MT prediction, the number of predictive dimensions is similar to the number of population dimensions, i.e., the predictive dimensions in MT are as large as possible and closely match the complexity of the target population, unlike the MT → SC prediction.

(B) Same as (A) for the attend out condition.

(C) Data in (A) and (B) presented as a ratio to compare the effect of attention on the communication subspace in MT.

(D) For the attend in condition, the number of private predictive dimensions is greater than the number of shared predictive dimensions in SC. For the SC → MT prediction (pink points), fewer dimensions are required to predict MT activity than are required to explain 95% of the variance in the MT activity, i.e., a communication subspace exists in SC that comprises ~2 shared dimensions that are sufficient to predict the ~3.5-dimensional activity in MT.

(E) Same as (D) for the attend out condition.

(F) Data in (D) and (E) presented as a ratio to compare the effect of attention on the communication subspace in SC.

principle be exclusive to pairs of regions that individually show significant changes in local representations. We tested this hypothesis by analyzing previously published simultaneous recordings of populations of neurons in V1 (which tend to show very modest effects of attention) and a single MT neuron.^{15,23,55} As with the MT → SC results, we found that attention dramatically improves V1 → MT prediction accuracy (Figure 7C; because we only recorded one MT neuron at a time, it was not possible to compute MT → V1 prediction accuracy). Also, the performance of V1-V1 predictions is not affected by attention (Figure 7D). These results demonstrate that, even though the effect of attention on V1 was small, attention-related effects on inter-areal communication are not contingent on large effects of attention in individual regions.

DISCUSSION

Our results show that attention changes the functional communication between populations of visual and premotor neurons. We demonstrated that attention changes the extent to which the activity of populations of neurons in the SC can be predicted by the activity of populations of neurons in MT and vice versa. These changes in functional interactions between areas are not accompanied by changes in the dimensionality of the subspace of activity that is shared between areas, and they are not obligatorily related to changes in firing rates, noise correlations, or population activity within each brain area. These results suggest that changes in information flow are well positioned to mediate behavioral flexibility and

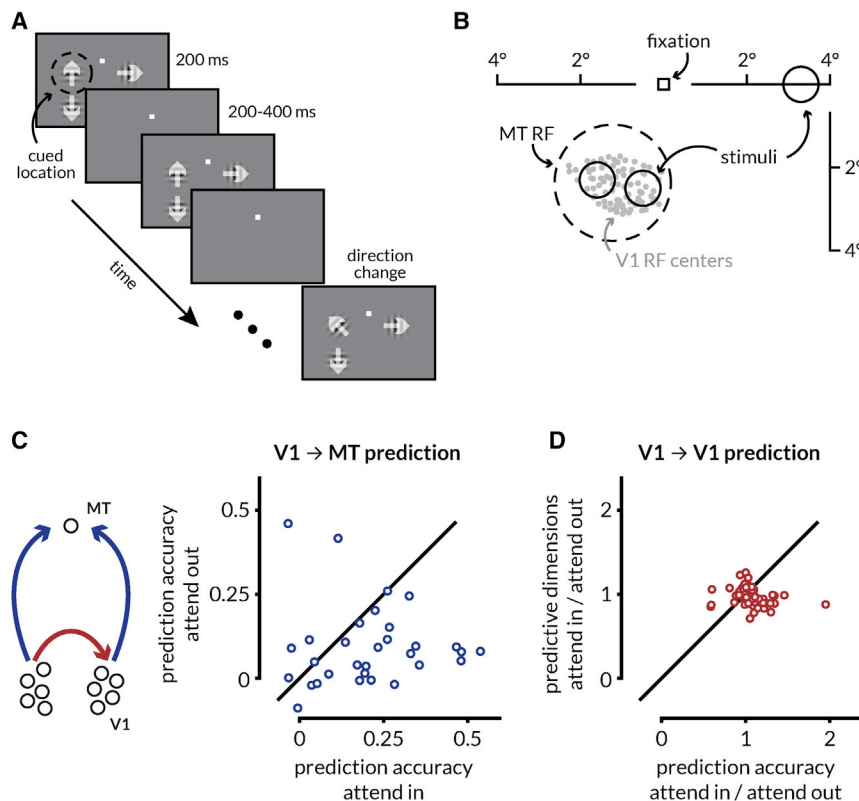


Figure 7. Attention enhances prediction accuracy between V1 and MT

(A) Schematic of the motion direction change detection task used during the V1-MT recordings. The monkeys were instructed to attend to changes in motion direction at one of three spatial locations while ignoring changes at the other two locations in blocks of 50–100 trials. The monkey started the trial by fixating a central spot. Two or three small Gabor stimuli synchronously flashed on for 200 ms and off for a randomized 200- to 400-ms period. Two of the stimuli were positioned inside the joint receptive fields (RFs) of the V1 and MT neurons, and the other was in the opposite hemifield. Trials during which attention was directed into the MT RF toward either of the two spatial locations were considered attend in trials, and trials in which attention was directed to the opposite hemifield were considered attend out trials. In blocks when the monkey was cued to attend to one of the two locations inside the RFs, the third stimulus was not presented. One of the two stimuli in the RF moved in the preferred direction of recorded MT neuron, and the other moved in the anti-preferred direction. When presented, the third stimulus moved in the preferred direction of the MT neuron. After a randomized number of stimulus presentations (between 2 and 13), the direction of one of the stimuli changed. The monkeys were rewarded for making a saccade to the direction change in the cued location. Premature saccades or saccades to changes in motion direction at the uncued location were not rewarded. We analyzed all identical stimulus presentations except the first to minimize the effect of adaptation.

(B) RF locations of recorded units from an example recording session. The gray dots represent the RF centers of 96 V1 neurons. The dotted circle represents the size and location of the RF for the recorded MT neuron. The size and locations of the stimuli were selected such that they lie within the MT RF.

(C) Attention improves the performance of V1 → MT prediction. Each dot represents the cross-validated performance for a linear model of the MT neuron's activity from V1 population activity using ridge regression for one recording session. The prediction accuracy on attend in trials was significantly greater than the accuracy on attend out trials ($p = 0.0159$; Wilcoxon signed-rank test). The value of the ridge parameter was chosen to be the smallest value for which the model performance was within 1 SEM of the peak performance.

(D) Attention does not affect the prediction accuracy or the number of predictive dimensions of intra-areal interactions. Each point represents the ratio of the average prediction performance and the average number of predictive dimensions of a linear model for random splits of the V1 population using RR regression.

place important constraints on models of flexible neural circuits.

How attention-related increases in functional communication fit in with hypothesized mechanisms underlying attention

Previous studies have focused on a small number of hypothesized mechanisms by which attention might improve perception.^{54,70–72} The most studied hypothesis is that attention improves perception by improving information encoding.^{6,11,16,73} The observed attention-related changes in the responses of individual neurons and in correlations between visual neurons appear consistent with this hypothesis. However, neuronal populations typically encode more than enough sensory information to account for psychophysical performance,^{14,17,49,74,75} and the changes in trial-by-trial fluctuations may not reflect changes in information coding that are behaviorally relevant.^{76,77} An alternate theory is that attention selectively improves the communication of sensory information to the neurons involved in perceptual decision making. Physiological studies along these lines have primarily focused on changes in synchrony or coherence between

areas on very short timescales (one or a few milliseconds; for review, see Womelsdorf and Fries⁴³) or using human imaging data to assess functional connectivity over multiple seconds.^{78–80} However, co-variability on short timescales is mathematically nearly independent of correlations on the timescale of hundreds of milliseconds,⁸¹ and unlike fluctuations on very short or very long timescales, response fluctuations on the timescale of hundreds of milliseconds covary with perceptual decisions.^{82,83}

Our results demonstrate that attention changes functional communication on the timescale of hundreds of milliseconds or seconds (Figure 5). This timescale can provide hints as to underlying mechanisms. For example, we showed previously that, in a network with spatially ordered, broad connectivity, modulating the balance of inhibition to excitation can change correlated variability on the timescale of hundreds of milliseconds.⁴⁵ Further, attention has been associated with changes in neuromodulators, including acetylcholine^{84,85} and dopamine,^{86–88} which operate on similar timescales as the attention-related changes in functional connectivity we observed.

Recently, we showed that attention is associated with only modest changes in either information coding in visual cortex or

the way information is read out by premotor neurons on the time-scale of perceptual decisions.¹⁴ Instead, our multi-neuron, multi-area recordings suggest that attention reshapes population activity in visual cortex, which changes the visual information that guides behavior via relatively fixed readout mechanisms. Our current results suggest a functional implication of this reshaping, changing the information that is shared between sensory neurons and the premotor neurons involved in decision making in a way that does not rely on obvious changes in the subspace of activity that is shared between them.

The communication subspace as a mechanism for flexible behavior

Several recent studies have shown that the activity of populations of neurons in many areas is generally confined to a subspace of population activity that is much lower dimensionally than the number of recorded neurons.^{54,56,89–101} The divergent anatomical connections between even the most highly interconnected brain areas have long suggested that only a portion of the information encoded in each area is shared between areas.

A recent set of studies used recordings from multiple populations of neurons to establish that functional communication between different brain areas in the motor⁹⁴ or visual system^{50,102} is confined to a subspace of activity that is even lower dimensionally than the activity within each area. Our results are consistent with the proposal that this limited communication subspace is an attractive mechanism for behavioral flexibility.^{50,94} Because only a subset of information is shared, reshaping activity within the source area (as in Ruff and Cohen¹⁴) and/or having a fixed but nonlinear subspace (proposed in Semedo et al.⁵⁰) would change the information that is functionally communicated to a target area. Using cross-prediction analyses, we found that these linear methods reveal a difference in the structure of the communication subspace across attention conditions, but this observation may be consistent with a fixed, non-linear communication subspace, and information flow could be improved by shifting the alignment of the shared fluctuations along the non-linearity (Figure S7). This mechanism is particularly attractive because changes in functional communication could occur without relying on changes in the weights relating one population to another, which may rely on synaptic plasticity mechanisms that occur over longer than behaviorally relevant timescales.⁴⁸

Our results suggest that the amount of information shared via the communication subspace between visual areas (V1 and MT; Figure 7) or between visual and premotor areas (MT and the SC; Figure 4) is in fact flexible. In future studies that explore a broader range of visual stimuli and behavioral tasks, it will be interesting to test the limits of this flexibility. For example, it will be important to know whether this mechanism might mediate flexible communication of different stimulus features or information accumulated on different timescales that must mediate more complex forms of behavioral flexibility.

An interesting aspect of spatial attention is that it is intimately linked to planning eye movements.^{103,104} This link is explicit in our task: when we cue a stimulus, we effectively instruct the monkey both to be better at the visual task of detecting a change and to plan an upcoming eye movement to that location. There is evidence that, even in tasks in which the behavioral response is

not an eye movement or when the eye movement is to a different location than the attended stimulus, attention also affects covert eye movement planning.¹⁰⁵ Our observation that attention affects functional interactions between V1 and MT, two areas with modest premotor signals, hints that the changes in functional communication we observed may not require changes in motor planning. In future work, it will be important to see whether attention can change functional communication using forms of attention that are not linked to motor planning (e.g., feature attention).^{106,107}

Constraints on mechanistic models

Measurements of the activity of large populations have proven critical for constraining mechanistic models. Phenomenological models can explain attention-related changes in firing rates,^{1,108–112} but these do not provide insight into circuit mechanisms. A staggering variety of biophysical models can recreate the effects of attention on the trial-averaged responses of individual neurons.^{45,113–118} We and others have shown that attention-related changes in correlated variability that resides in a low dimensional subspace of population activity provides much stronger constraints on circuit models.^{45,116}

The observation that functional communication between areas is lower dimensional compared to activity within each area^{51,94} and our observation that attention changes this communication will further constrain circuit models. In particular, many models^{45,116,119–121} and experiments^{122–124} implicate inhibition in the flexibility of neuronal populations, but early efforts suggest that these mechanisms do not readily create low dimensional and flexible communication subspaces (Gozel and Doiron, 2021, Cosyne, poster). It is possible that the complementary influence of different subtypes of inhibitory interneurons may underlie the flexible functional communication we observed.^{84,125–127} Using observations about interactions between neural populations to constrain mechanistic, generative models will be critical for understanding how neural circuits give rise to behavior in a wide variety of systems, species, and disease states.

Concluding remarks

The hallmark of the nervous system is its flexibility. Flexible behavior must rely, on some level, on flexible information flow. Attention, which changes the behavioral importance of different objects, features, or locations, is a good model of flexible information flow. Our results demonstrate that this flexibility is instantiated, at least in part, by changes in the information that is shared between different stages of the visuomotor pathway. These results lay the groundwork for establishing the role of flexible inter-area communications in a variety of sensory, cognitive, and motor computations.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead contact

- Materials availability
- Data and code availability
- **EXPERIMENTAL MODEL AND SUBJECT DETAILS**
- **METHOD DETAILS**
 - Electrophysiological Recordings and Behavioral Task
 - MT-SC recordings
 - V1-MT recordings
- **QUANTIFICATION AND STATISTICAL ANALYSIS**
 - Subsampling
 - Noise correlations
 - Regression
 - Long timescale correlations and regression
 - Cross-condition, cross-validated regression
 - Factor Analysis

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.cub.2021.09.076>.

ACKNOWLEDGMENTS

We are grateful to K. McRackeen for providing technical assistance, to Adam Kohn for comments on an earlier version of this manuscript, and to Brent Doiron, Joao Smedo, and Byron Yu for helpful comments and suggestions regarding data analysis. M.R.C. is supported by National Institutes of Health grants R01EY022930 and R01NS121913, and Simons Foundation grant 542961SPI.

AUTHOR CONTRIBUTIONS

Conceptualization, methodology, writing – review & editing, R.S., D.A.R., and M.R.C.; software, R.S. and D.A.R.; analysis, visualization, writing – original draft, R.S. and M.R.C.; funding acquisition, resources, supervision, M.R.C.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: April 29, 2021

Revised: September 22, 2021

Accepted: September 28, 2021

Published: October 25, 2021

REFERENCES

1. Maunsell, J.H.R. (2015). Neuronal mechanisms of visual attention. *Annu. Rev. Vis. Sci.* 1, 373–391.
2. Moore, T., and Zirnsak, M. (2017). Neural mechanisms of selective visual attention. *Annu. Rev. Psychol.* 68, 47–72.
3. Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222.
4. Reynolds, J.H., and Chelazzi, L. (2004). Attentional modulation of visual processing. *Annu. Rev. Neurosci.* 27, 611–647.
5. Cohen, M.R., and Maunsell, J.H.R. (2011). Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron* 70, 1192–1204.
6. Cohen, M.R., and Maunsell, J.H.R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* 12, 1594–1600.
7. Gregoriou, G.G., Rossi, A.F., Ungerleider, L.G., and Desimone, R. (2014). Lesions of prefrontal cortex reduce attentional modulation of neuronal responses and synchrony in V4. *Nat. Neurosci.* 17, 1003–1011.
8. Herrero, J.L., Gieselmann, M.A., Sanayei, M., and Thiele, A. (2013). Attention-induced variance and noise correlation reduction in macaque V1 is mediated by NMDA receptors. *Neuron* 78, 729–739.
9. Luo, T.Z., and Maunsell, J.H.R. (2015). Neuronal modulations in visual cortex are associated with only one of multiple components of attention. *Neuron* 86, 1182–1188.
10. Mayo, J.P., and Maunsell, J.H.R. (2016). Graded neuronal modulations related to visual spatial attention. *J. Neurosci.* 36, 5353–5361.
11. Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63, 879–888.
12. Nandy, A.S., Nassi, J.J., and Reynolds, J.H. (2017). Laminar organization of attentional modulation in macaque visual area V4. *Neuron* 93, 235–246.
13. Ni, A.M., Ruff, D.A., Alberts, J.J., Symmonds, J., and Cohen, M.R. (2018). Learning and attention reveal a general relationship between population activity and behavior. *Science* 359, 463–465.
14. Ruff, D.A., and Cohen, M.R. (2019). Simultaneous multi-area recordings suggest that attention improves performance by reshaping stimulus representations. *Nat. Neurosci.* 22, 1669–1676.
15. Ruff, D.A., and Cohen, M.R. (2016). Attention increases spike count correlations between visual cortical areas. *J. Neurosci.* 36, 7523–7534.
16. Ruff, D.A., and Cohen, M.R. (2014). Attention can either increase or decrease spike count correlations in visual cortex. *Nat. Neurosci.* 17, 1591–1597.
17. Ruff, D.A., and Cohen, M.R. (2014). Global cognitive factors modulate correlated response variability between V4 neurons. *J. Neurosci.* 34, 16408–16416.
18. Verhoef, B.-E., and Maunsell, J.H.R. (2017). Attention-related changes in correlated neuronal activity arise from normalization mechanisms. *Nat. Neurosci.* 20, 969–977.
19. Yan, Y., Rasch, M.J., Chen, M., Xiang, X., Huang, M., Wu, S., and Li, W. (2014). Perceptual training continuously refines neuronal population codes in primary visual cortex. *Nat. Neurosci.* 17, 1380–1387.
20. Zénon, A., and Krauzlis, R.J. (2012). Attention deficits without cortical neuronal deficits. *Nature* 489, 434–437.
21. Oemisch, M., Westendorff, S., Everling, S., and Womelsdorf, T. (2015). Interareal spike-train correlations of anterior cingulate and dorsal prefrontal cortex during attention shifts. *J. Neurosci.* 35, 13076–13089.
22. Pooremailli, A., Poort, J., and Roelfsema, P.R. (2014). Simultaneous selection by object-based attention in visual and frontal cortex. *Proc. Natl. Acad. Sci. USA* 111, 6467–6472.
23. Ruff, D.A., Alberts, J.J., and Cohen, M.R. (2016). Relating normalization to neuronal populations across cortical areas. *J. Neurophysiol.* 116, 1375–1386.
24. Azouz, R., and Gray, C.M. (2003). Adaptive coincidence detection and dynamic gain control in visual cortical neurons in vivo. *Neuron* 37, 513–523.
25. Bichot, N.P., Rossi, A.F., and Desimone, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* 308, 529–534.
26. Bosman, C.A., Schoffelen, J.-M., Brunet, N., Oostenveld, R., Bastos, A.M., Womelsdorf, T., Rubehn, B., Stieglitz, T., De Weerd, P., and Fries, P. (2012). Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron* 75, 875–888.
27. Briggs, F., Mangun, G.R., and Usrey, W.M. (2013). Attention enhances synaptic efficacy and the signal-to-noise ratio in neural circuits. *Nature* 499, 476–480.
28. Buffalo, E.A., Fries, P., Landman, R., Buschman, T.J., and Desimone, R. (2011). Laminar differences in gamma and alpha coherence in the ventral stream. *Proc. Natl. Acad. Sci. USA* 108, 11262–11267.

29. Buschman, T.J., and Miller, E.K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* 315, 1860–1862.
30. Dagnino, B., Gariel-Mathis, M.-A., and Roelfsema, P.R. (2015). Microstimulation of area V4 has little effect on spatial attention and on perception of phosphenes evoked in area V1. *J. Neurophysiol.* 113, 730–739.
31. Fries, P. (2015). Rhythms for cognition: communication through coherence. *Neuron* 88, 220–235.
32. Fries, P., Reynolds, J.H., Rorie, A.E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291, 1560–1563.
33. Gregoriou, G.G., Gotts, S.J., Zhou, H., and Desimone, R. (2009). High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* 324, 1207–1210.
34. Klink, P.C., Dagnino, B., Gariel-Mathis, M.-A., and Roelfsema, P.R. (2017). Distinct feedforward and feedback effects of microstimulation in visual cortex reveal neural mechanisms of texture segregation. *Neuron* 95, 209–220.e3.
35. Lakatos, P., Karmos, G., Mehta, A.D., Ulbert, I., and Schroeder, C.E. (2008). Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* 320, 110–113.
36. Miller, E.K., and Buschman, T.J. (2013). Cortical circuits for the control of attention. *Curr. Opin. Neurobiol.* 23, 216–222.
37. Moore, T., and Armstrong, K.M. (2003). Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 421, 370–373.
38. Ruff, D.A., and Cohen, M.R. (2017). A normalization model suggests that attention changes the weighting of inputs between visual areas. *Proc. Natl. Acad. Sci. USA* 114, E4085–E4094.
39. Salmann, Y.B., Pigarev, I.N., and Vidyasagar, T.R. (2007). Neural mechanisms of visual attention: how top-down feedback highlights relevant locations. *Science* 316, 1612–1615.
40. Salinas, E., and Sejnowski, T.J. (2001). Correlated neuronal activity and the flow of neural information. *Nat. Rev. Neurosci.* 2, 539–550.
41. Saproo, S., and Serences, J.T. (2014). Attention improves transfer of motion information between V1 and MT. *J. Neurosci.* 34, 3586–3596.
42. Womelsdorf, T., Fries, P., Mitra, P.P., and Desimone, R. (2006). Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* 439, 733–736.
43. Womelsdorf, T., and Fries, P. (2007). The role of neuronal synchronization in selective attention. *Curr. Opin. Neurobiol.* 17, 154–160.
44. Cowley, B.R., Snyder, A.C., Acar, K., Williamson, R.C., Yu, B.M., and Smith, M.A. (2020). Slow drift of neural activity as a signature of impulsivity in macaque visual and prefrontal cortex. *Neuron* 108, 551–567.e8.
45. Huang, C., Ruff, D.A., Pyle, R., Rosenbaum, R., Cohen, M.R., and Doiron, B. (2019). Circuit models of low-dimensional shared variability in cortical networks. *Neuron* 101, 337–348.e4.
46. Ruff, D.A., Xue, C., Kramer, L.E., Baqai, F., and Cohen, M.R. (2020). Low rank mechanisms underlying flexible visual representations. *Proc. Natl. Acad. Sci. USA* 117, 29321–29329.
47. Carrasco, M. (2011). Visual attention: the past 25 years. *Vision Res.* 51, 1484–1525.
48. Egeth, H.E., and Yantis, S. (1997). Visual attention: control, representation, and time course. *Annu. Rev. Psychol.* 48, 269–297.
49. Kohn, A., Coen-Cagli, R., Kanitscheider, I., and Pouget, A. (2016). Correlations and neuronal population information. *Annu. Rev. Neurosci.* 39, 237–256.
50. Smedo, J.D., Zandvakili, A., Machens, C.K., Yu, B.M., and Kohn, A. (2019). Cortical areas interact through a communication subspace. *Neuron* 102, 249–259.e4.
51. Smedo, J.D., Gokcen, E., Machens, C.K., Kohn, A., and Yu, B.M. (2020). Statistical methods for dissecting interactions between brain areas. *Curr. Opin. Neurobiol.* 65, 59–69.
52. Cohen, M.R., and Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nat. Neurosci.* 14, 811–819.
53. Umakantha, A., Morina, R., Cowley, B.R., Snyder, A.C., Smith, M.A., and Yu, B.M. (2021). Bridging neuronal correlations and dimensionality reduction. *Neuron* 109, 2740–2754.e12.
54. Ruff, D.A., Ni, A.M., and Cohen, M.R. (2018). Cognition as a window into neuronal population space. *Annu. Rev. Neurosci.* 41, 77–97.
55. Ruff, D.A., and Cohen, M.R. (2016). Stimulus dependence of correlated variability across cortical areas. *J. Neurosci.* 36, 7546–7556.
56. Cunningham, J.P., and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* 17, 1500–1509.
57. Fries, W. (1985). Inputs from motor and premotor cortex to the superior colliculus of the macaque monkey. *Behav. Brain Res.* 18, 95–105.
58. Fries, W. (1984). Cortical projections to the superior colliculus in the macaque monkey: a retrograde study using horseradish peroxidase. *J. Comp. Neurol.* 230, 55–76.
59. Lock, T.M., Baizer, J.S., and Bender, D.B. (2003). Distribution of corticotectal cells in macaque. *Exp. Brain Res.* 151, 455–470.
60. Lyon, D.C., Nassi, J.J., and Callaway, E.M. (2010). A disinaptic relay from superior colliculus to dorsal stream visual cortex in macaque monkey. *Neuron* 65, 270–279.
61. Stepniewska, I., Qi, H.-X., and Kaas, J.H. (1999). Do superior colliculus projection zones in the inferior pulvinar project to MT in primates? *Eur. J. Neurosci.* 11, 469–480.
62. Benevento, L.A., and Yoshida, K. (1981). The afferent and efferent organization of the lateral geniculo-prestriate pathways in the macaque monkey. *J. Comp. Neurol.* 203, 455–474.
63. Rodman, H.R., Gross, C.G., and Albright, T.D. (1990). Afferent basis of visual response properties in area MT of the macaque. II. Effects of superior colliculus removal. *J. Neurosci.* 10, 1154–1164.
64. Goldberg, M.E., and Wurtz, R.H. (1972). Activity of superior colliculus in behaving monkey. II. Effect of attention on neuronal responses. *J. Neurophysiol.* 35, 560–574.
65. Ignashchenkova, A., Dicke, P.W., Haarmeier, T., and Thier, P. (2004). Neuron-specific contribution of the superior colliculus to overt and covert shifts of attention. *Nat. Neurosci.* 7, 56–64.
66. Krauzlis, R.J., Lovejoy, L.P., and Zénon, A. (2013). Superior colliculus and visual spatial attention. *Annu. Rev. Neurosci.* 36, 165–182.
67. Recanzone, G.H., and Wurtz, R.H. (2000). Effects of attention on MT and MST neuronal activity during pursuit initiation. *J. Neurophysiol.* 83, 777–790.
68. Seidemann, E., and Newsome, W.T. (1999). Effect of spatial attention on the responses of area MT neurons. *J. Neurophysiol.* 81, 1783–1794.
69. Womelsdorf, T., Anton-Erxleben, K., Pieper, F., and Treue, S. (2006). Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. *Nat. Neurosci.* 9, 1156–1160.
70. Driver, J. (2001). A selective review of selective attention research from the past century. *Br. J. Psychol.* 92, 53–78.
71. Lavie, N. (2010). Attention, distraction, and cognitive control under load. *Curr. Dir. Psychol. Sci.* 19, 143–148.
72. Peelen, M.V., and Kastner, S. (2014). Attention in the real world: toward understanding its neural basis. *Trends Cogn. Sci.* 18, 242–250.
73. Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* 55, 131–141.
74. Kanitscheider, I., Coen-Cagli, R., and Pouget, A. (2015). Origin of information-limiting noise correlations. *Proc. Natl. Acad. Sci. USA* 112, E6973–E6982.
75. Parker, A.J., and Newsome, W.T. (1998). Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci.* 21, 227–277.
76. Baruni, J.K., Lau, B., and Salzman, C.D. (2015). Reward expectation differentially modulates attentional behavior and activity in visual area V4. *Nat. Neurosci.* 18, 1656–1663.

77. Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., and Pouget, A. (2014). Information-limiting correlations. *Nat. Neurosci.* **17**, 1410–1417.
78. Indovina, I., and Macaluso, E. (2004). Occipital-parietal interactions during shifts of exogenous visuospatial attention: trial-dependent changes of effective connectivity. *Magn. Reson. Imaging* **22**, 1477–1486.
79. Ozaki, T.J. (2011). Frontal-to-parietal top-down causal streams along the dorsal attention network exclusively mediate voluntary orienting of attention. *PLoS ONE* **6**, e20079.
80. Rossi, S., Huang, S., Furtak, S.C., Belliveau, J.W., and Ahveninen, J. (2014). Functional connectivity of dorsal and ventral frontoparietal seed regions during auditory orienting. *Brain Res.* **1583**, 159–168.
81. Bair, W., Zohary, E., and Newsome, W.T. (2001). Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J. Neurosci.* **21**, 1676–1697.
82. Nienborg, H., Cohen, M.R., and Cumming, B.G. (2012). Decision-related activity in sensory neurons: correlations among neurons and with behavior. *Annu. Rev. Neurosci.* **35**, 463–483.
83. Nienborg, H., and Cumming, B. (2010). Correlations between the activity of sensory neurons and behavior: how much do they tell us about a neuron's causality? *Curr. Opin. Neurobiol.* **20**, 376–381.
84. Herrero, J.L., Roberts, M.J., Delicato, L.S., Gieselmann, M.A., Dayan, P., and Thiele, A. (2008). Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* **454**, 1110–1114.
85. Thiele, A., and Bellgrove, M.A. (2018). Neuromodulation of attention. *Neuron* **97**, 769–785.
86. Noudoost, B., and Moore, T. (2011). Control of visual cortical signals by prefrontal dopamine. *Nature* **474**, 372–375.
87. Soltani, A., Noudoost, B., and Moore, T. (2013). Dissociable dopaminergic control of saccadic target selection and its implications for reward modulation. *Proc. Natl. Acad. Sci. USA* **110**, 3579–3584.
88. Noudoost, B., and Moore, T. (2011). The role of neuromodulators in selective attention. *Trends Cogn. Sci.* **15**, 585–591.
89. Cowley, B.R., Smith, M.A., Kohn, A., and Yu, B.M. (2016). Stimulus-driven population activity patterns in macaque primary visual cortex. *PLoS Comput. Biol.* **12**, e1005185.
90. Elsayed, G.F., Lara, A.H., Kaufman, M.T., Churchland, M.M., and Cunningham, J.P. (2016). Reorganization between preparatory and movement population responses in motor cortex. *Nat. Commun.* **7**, 13239.
91. Elsayed, G.F., and Cunningham, J.P. (2017). Structure in neural population recordings: an expected byproduct of simpler phenomena? *Nat. Neurosci.* **20**, 1310–1318.
92. Golub, M.D., Chase, S.M., Batista, A.P., and Yu, B.M. (2016). Brain-computer interfaces for dissecting cognitive processes underlying sensorimotor control. *Curr. Opin. Neurobiol.* **37**, 53–58.
93. Jazayeri, M., and Afraz, A. (2017). Navigating the neural space in search of the neural code. *Neuron* **93**, 1003–1014.
94. Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* **17**, 440–448.
95. Kiani, R., Esteky, H., Mirpour, K., and Tanaka, K. (2007). Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *J. Neurophysiol.* **97**, 4296–4309.
96. Miri, A., Warriner, C.L., Seely, J.S., Elsayed, G.F., Cunningham, J.P., Churchland, M.M., and Jessell, T.M. (2017). Behaviorally selective engagement of short-latency effector pathways by motor cortex. *Neuron* **95**, 683–696.e11.
97. Morcos, A.S., and Harvey, C.D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. *Nat. Neurosci.* **19**, 1672–1681.
98. Pandarinath, C., O'Shea, D.J., Collins, J., Jozefowicz, R., Stavisky, S.D., Kao, J.C., Trautmann, E.M., Kaufman, M.T., Ryu, S.I., Hochberg, L.R., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat. Methods* **15**, 805–815.
99. Pitkow, X., and Angelaki, D.E. (2017). Inference in the brain: statistics flowing in redundant population codes. *Neuron* **94**, 943–953.
100. Sadtler, P.T., Quick, K.M., Golub, M.D., Chase, S.M., Ryu, S.I., Tyler-Kabara, E.C., Yu, B.M., and Batista, A.P. (2014). Neural constraints on learning. *Nature* **512**, 423–426.
101. Yu, B.M., Cunningham, J.P., Santhanam, G., Ryu, S.I., Shenoy, K.V., and Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* **102**, 614–635.
102. Smedo, J.D., Jasper, A.I., Zandvakili, A., Aschner, A., Machens, C.K., Kohn, A., and Yu, B.M. (2021). Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *bioRxiv*. <https://doi.org/10.1101/2021.02.08.430346>.
103. Bisley, J.W., and Goldberg, M.E. (2010). Attention, intention, and priority in the parietal lobe. *Annu. Rev. Neurosci.* **33**, 1–21.
104. Andersen, R.A., and Buneo, C.A. (2002). Intentional maps in posterior parietal cortex. *Annu. Rev. Neurosci.* **25**, 189–220.
105. Steinmetz, N.A., and Moore, T. (2014). Eye movement preparation modulates neuronal responses in area V4 when dissociated from attentional demands. *Neuron* **83**, 496–506.
106. Snyder, L.H., Batista, A.P., and Andersen, R.A. (1997). Coding of intention in the posterior parietal cortex. *Nature* **386**, 167–170.
107. Maunsell, J.H.R., and Treue, S. (2006). Feature-based attention in visual cortex. *Trends Neurosci.* **29**, 317–322.
108. Boynton, G.M. (2009). A framework for describing the effects of attention on visual responses. *Vision Res.* **49**, 1129–1143.
109. Ecker, A.S., Denfield, G.H., Bethge, M., and Tolias, A.S. (2016). On the structure of neuronal population activity under fluctuations in attentional state. *J. Neurosci.* **36**, 1775–1789.
110. Gilbert, C.D., and Sigman, M. (2007). Brain states: top-down influences in sensory processing. *Neuron* **54**, 677–696.
111. Navalpakkam, V., and Itti, L. (2005). Modeling the influence of task on attention. *Vision Res.* **45**, 205–231.
112. Reynolds, J.H., and Heeger, D.J. (2009). The normalization model of attention. *Neuron* **61**, 168–185.
113. Ardid, S., Wang, X.-J., and Compte, A. (2007). An integrated microcircuit model of attentional processing in the neocortex. *J. Neurosci.* **27**, 8486–8495.
114. Buia, C.I., and Tiesinga, P.H. (2008). Role of interneuron diversity in the cortical microcircuit for attention. *J. Neurophysiol.* **99**, 2158–2182.
115. Deco, G., and Thiele, A. (2011). Cholinergic control of cortical network interactions enables feedback-mediated attentional modulation. *Eur. J. Neurosci.* **34**, 146–157.
116. Kanashiro, T., Ocker, G.K., Cohen, M.R., and Doiron, B. (2017). Attentional modulation of neuronal variability in circuit models of cortex. *eLife* **6**, e23978.
117. Silver, R.A. (2010). Neuronal arithmetic. *Nat. Rev. Neurosci.* **11**, 474–489.
118. Sutherland, M.R., McQuiggan, D.A., Ryan, J.D., and Mather, M. (2017). Perceptual salience does not influence emotional arousal's impairing effects on top-down attention. *Emotion* **17**, 700–706.
119. Brunel, N., and Wang, X.-J. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J. Comput. Neurosci.* **11**, 63–85.
120. Machens, C.K., Romo, R., and Brody, C.D. (2005). Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* **307**, 1121–1124.
121. Rubin, D.B., Van Hooser, S.D., and Miller, K.D. (2015). The stabilized supralinear network: a unifying circuit motif underlying multi-input integration in sensory cortex. *Neuron* **85**, 402–417.

122. Fu, Y., Tucciarone, J.M., Espinosa, J.S., Sheng, N., Darcy, D.P., Nicoll, R.A., Huang, Z.J., and Stryker, M.P. (2014). A cortical circuit for gain control by behavioral state. *Cell* 156, 1139–1152.
123. Karnani, M.M., Jackson, J., Ayzenshtat, I., Hamzehei Sichani, A., Manoocheri, K., Kim, S., and Yuste, R. (2016). Opening holes in the blanket of inhibition: localized lateral disinhibition by VIP interneurons. *J. Neurosci.* 36, 3471–3480.
124. Kuchibhotla, K.V., Gill, J.V., Lindsay, G.W., Papadoyannis, E.S., Field, R.E., Sten, T.A.H., Miller, K.D., and Froemke, R.C. (2017). Parallel processing by cortical inhibition enables context-dependent behavior. *Nat. Neurosci.* 20, 62–71.
125. Cardin, J.A., Carlén, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Tsai, L.-H., and Moore, C.I. (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature* 459, 663–667.
126. Roberts, M.J., Zinke, W., Guo, K., Robertson, R., McDonald, J.S., and Thiele, A. (2005). Acetylcholine dynamically controls spatial integration in marmoset primary visual cortex. *J. Neurophysiol.* 93, 2062–2072.
127. Veit, J., Hakim, R., Jädi, M.P., Sejnowski, T.J., and Adesnik, H. (2017). Cortical gamma band synchronization through somatostatin interneurons. *Nat. Neurosci.* 20, 951–959.
128. Brainard, D.H. (1997). The Psychophysics Toolbox. *Spat. Vis.* 10, 433–436.
129. Everitt, B.S. (1984). Maximum likelihood estimation of the parameters in a mixture of two univariate normal distributions; a comparison of different algorithms. *J. R. Stat. Soc. Ser. Stat.* 33, 205–215.
130. Semedo, J.D., Zandvakili, A., Kohn, A., Machens, C.K., and Yu, B.M. (2014). Extracting latent structure from multiple interacting neural populations. In *NIPS'14: Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2* (MIT), pp. 2942–2950.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental models: Organisms/strains		
Rhesus Macaques (<i>Macaca mulatta</i>)	University of Pittsburgh and Carnegie Mellon University	N/A
Software and algorithms		
MATLAB	MathWorks	https://www.mathworks.com/products/MATLAB.html
Psychophysics Toolbox v3	Brainard ¹²⁸	http://psychtoolbox.org
Other		
Plexon 24-channel linear probes	Plexon: Dallas, TX 75206 USA	https://plexon.com/products/plexon-s-probe/
Blackrock 10x10 array	Blackrock Microsystems: Salt Lake City, UT 84108 USA	http://blackrockmicro.com
Ripple Neuromed	Ripple Neuromed: Salt Lake City, UT 84106 USA	http://rippleneckneuromed.com
EyeLink 1000 Eye tracking system	SR Research: Ottawa, Ontario, Canada., K2L 2B9	http://sr-research.com
Plexon Offline Sorter 3.3.2	Plexon: Dallas, TX 75206 USA	https://plexon.com/products/offline-sorter

RESOURCE AVAILABILITY

Lead contact

Requests for resources should be directed to and will be fulfilled by the Lead Contact, Ramanujan Srinath (ramanujan@pitt.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

The data and code that support the findings of this study have been deposited in a public Github repository <https://github.com/ramanujansrinath/mt-sc-comm-data>. MATLAB code for reduced-rank regression and factor analysis has been made publicly available by Byron Yu and can be downloaded from <https://users.ece.cmu.edu/~byronyu/software.shtml>. Further information and requests for data or custom MATLAB code should be directed to and will be fulfilled by the Lead Contact, Ramanujan Srinath (ramanujan@pitt.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

The electrophysiological data in this manuscript comes from two previously reported experiments.^{14,15} In both experiments, two adult male rhesus monkeys (*Macaca mulatta*, 8 and 9 kg) were used. We implanted each animal with a titanium head post before behavioral training. We identified each cortical area by visualizing the sulci during array implantation, using stereotactic coordinates, and by observing the transition of gray and white matter signals on the movable probes. All animal procedures were approved by the Institutional Animal Care and Use Committees of the University of Pittsburgh and Carnegie Mellon University.

METHOD DETAILS

Electrophysiological Recordings and Behavioral Task

Our methods for presenting visual stimuli and monitoring behavior have been described previously. Briefly, we presented visual stimuli using custom software (written in MATLAB using the Psychophysics Toolbox v3¹²⁸ on a cathode-ray tube monitor (calibrated to linearize intensity; 1,024 × 768 pixels; 120 Hz refresh rate) placed 54 cm from each animal. We monitored eye position using an infrared eye tracker (EyeLink 1000; SR Research) and recorded eye position and pupil diameter (1,000 samples/s), neuronal responses (30,000 samples/s) and the signal from a photodiode to align neuronal responses to stimulus presentation times (30,000 samples/s) using hardware from Ripple. All spike sorting was done offline manually using Offline Sorter (version 3.3.2; Plexon). We based our analyses on both single units and multi-unit clusters and use the term *unit* to refer to either.

In both sets of recordings described below, monkeys were trained to do a stimulus direction change detection task. Monkeys were cued to attend to a spatial location (typically within or away from the receptive fields of the neurons). While the monkeys fixated on a central spot, identical stimuli were flashed at two/three locations on the screen and, at an unsignaled time, one of them would change direction. The monkey was rewarded for making saccade toward the changed stimulus. Further details regarding stimulus locations, timing, etc. are specific to the two recording experiments and are described below.

MT-SC recordings

We implanted two recording chambers on the right hemisphere that granted access to MT and SC for recordings with linear 24-channel moveable probes (Plexon; interelectrode spacing in MT = 50 μm , SC = 100 μm) and simultaneously recorded activity from neurons in MT and SC that had overlapping spatial receptive fields (Figure 1). To account for visual latencies in the two areas, spikes were counted between 50 and 250ms after stimulus onset. We only analyzed a recorded MT unit if its stimulus-driven firing rate was 10% higher than its firing rate as measured in the 100ms before 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 100ms before the onset of the first stimulus or if its response during a 100ms epoch before a saccade on hit (correct) trials to the contralateral side was 10% larger than that same baseline. The dataset consisted of a total of 306 responsive MT units (6–29 units per session, mean 20.4) and 345 responsive SC units (12–29 units per session, mean 23) across 15 recording sessions (6 and 9 sessions from monkey 1 and 2 respectively). Approximately 11.5% and 9.6% of MT and SC units respectively were single units isolated post hoc using an offline spike sorting program – MT: 0–8 (mean 2.2) single units, 5–24 (mean 18.1) multi-units per session; SC: 0–6 (mean 2.2) single units, 12–24 (mean 20.8) multi-units per session.

Experimental task design

Each session began with receptive field mapping using a delayed-saccade task, and direction tuning during passive fixation, followed by multiple blocks of the following attention task. Each block began with a set of trials that instructed the monkey to attend to one of two spatial locations on the screen – either within the joint receptive fields of the neurons or in the opposite hemifield. Following that, each trial began when the monkey acquired fixation on a central spot within a 1.25° fixation window after which two peripheral drifting Gabor stimuli (one overlapping the receptive fields of the recorded neurons, the other in the opposite visual hemifield) synchronously flashed on (for 200ms) and off (for a randomized period between 200 and 400ms) between 3–12 times before, at a random, unsignaled time, the direction of one of the stimuli changed from that of the preceding stimulus. The monkey reported the orientation change by making a saccade to the changed stimulus within 450ms and received a juice reward. Each block consisted of approximately 100 completed trials (i.e., trials that ended in a hit or miss) after which the cued location of the orientation change switched to the other hemifield. Stimulus presentations during the response period of which the monkey made a micro-saccade were excluded from analysis. Neural responses to all stimulus presentations after the first (to minimize the effect of adaptation) and before the orientation change were analyzed. For each session, stimulus presentations were sampled such that the number of presentations was equal for each attention condition. Each session yielded 547–1909 (mean 1277) presentations for each attention condition. The attentional modulation index for each neuron was calculated as a difference of mean responses between the two attention conditions divided by the sum.

For each session, SC neurons were divided evenly into oculo-motor (*visual* for brevity) and motor neurons based on an oculo-motor score calculated as

$$\text{score}_{\text{vis/mot}} = (R_{\text{vis}} - R_{\text{mot}}) / (R_{\text{vis}} + R_{\text{mot}})$$

where R_{vis} is the average neural response to the onset of the first stimulus, and R_{mot} is the average response prior to a saccade to the stimulus in the contralateral hemifield. This score was calculated for the trials where attention was directed into the joint RFs.

V1-MT recordings

We implanted a 10x10 chronic microelectrode array (Blackrock Microsystems) in V1 and a recording chamber to access MT. Each recording session began with searching a well-isolated MT neuron such that its receptive field (RF) overlapped the population RF of the V1 neurons and was driven similarly above baseline by a single stimulus flashed in each of two chosen locations. This dataset consisted of a total of 1631 responsive V1 units and 32 responsive MT units (1 unit per session in MT, 7–83 units per session, mean 51 in V1) across 32 recording sessions.

Experimental task design

Each block of trials began with a set of trials that instructed the monkey to attend to one of three spatial locations on the screen – either one of two locations within the receptive field of the MT neuron or one in the opposite hemifield. Each trial began when the monkey acquired fixation on a 1° fixation window. For blocks in which attention was directed within the RF of the MT neuron, two achromatic Gabor stimuli of equal contrast, spatial frequency, and speed were presented drifting in opposite directions (preferred and null direction for the MT neuron). For blocks in which attention was directed to the opposite hemifield, a third drifting Gabor was similarly flashed at the cued location. In these blocks, the contrast of the stimulus at the cued location was different from the two stimuli within the RF of the MT neuron. This was done to study the stimulus dependence of spike count correlations across cortical areas but is not critical to the current analyses as here the comparison is between the trials where attention is directed either into or out of the RF of the MT neuron regardless of stimulus parameters or specific location with the RF. After 2–14 presentations of

the same stimuli, the direction of the stimulus at the cued location was changed and the monkey was rewarded for making a saccade to the changed stimulus within 500ms. As with the MT-SC data, stimulus presentations during which the monkey made a micro-saccade were excluded from analysis, all stimulus presentations after the first and before the orientation change were analyzed, and the presentations were sampled such that they were equal in the two attention conditions. Each session yielded 97–1469 (mean 583) presentations for each attention condition.

QUANTIFICATION AND STATISTICAL ANALYSIS

Subsampling

To test whether attention affects prediction of neural responses within and across areas, we first sought to check whether or not the number of recorded neurons and trials across the two attention conditions in the datasets were sufficient for reasonable regression performance. We used a linear model of the form $Y = XB$ where X and Y are matrices of $t \times n$ and $t \times m$ dimensions and B is the weight matrix of dimensions $n \times m$ (here, t is the number of stimulus presentations in a session, m and n are the numbers of neurons in the two areas). We found the ordinary least-squares solution for B by minimizing the squared prediction error as $B = (X^T X)^{-1} X^T Y$. We sampled N MT neurons (where N went from 1 to the total number of recorded MT neurons) without replacement and used ridge regression to predict SC responses. We did this subsampling 100 times for each N . For ridge regression, we chose the value of the regularization parameter (λ) using 10-fold cross-validation. The reported prediction accuracy is for the model with the largest λ for which mean performance (across folds) was within one SEM of the best performance across models. We also used the full MT recorded population to predict the responses of subsets of N SC neurons (where N went from 1 to the total number of recorded SC neurons) using the same method.

Noise correlations

The spike count correlation (r_{SC}) was calculated as the correlation coefficient between the responses of the two units to repeated presentations of the same stimulus. Z-scoring responses before calculating noise correlations did not qualitatively change the comparisons between noise correlations and local or shared dimensionality or prediction accuracy. In [Figure S1](#), noise correlations are computed for each pair in a session using all stimulus presentations in every trial (except the first), and then pooled across sessions and monkeys to yield 3315 pairs in MT, 3975 pairs in SC, and 6934 pairs across MT and SC. In [Figure S3](#), noise correlations are computed as above and then averaged for each session.

Regression

To find the effect of attention on the ability of MT responses to predict SC responses and vice versa, we used the same linear model described above using ridge regression. This model is referred to as the full regression model in the text. To assess whether the SC activity can be predicted using a subset of MT population response dimensions (in other words, a subspace of MT activity), we used reduced-rank (RR) regression. The exact description and formulation of RR regression can be found in Semedo et al.⁵⁰ Briefly, RR regression constrains the weight matrix B to be of a given rank and is solved using singular value decomposition:

$$Y_{RRR} = XB_{RRR} = XB_{OLS}VV^T = XB_{RRR}^T$$

where B_{OLS} is the coefficient matrix for the ordinary least-square solution, B_{RRR} is the coefficient matrix for the RR regression solution, V is a matrix whose columns contain the top principal components of the optimal linear predictor $Y_{OLS} = XB_{OLS}$. The columns of B define which dimensions of X are used for generating the prediction i.e., the predictive dimensions. As with the ridge regression solution above, we used 10-fold cross-validation and found the smallest number of dimensions for which predictive performance was within one SEM of the peak performance. The prediction performance of both ridge and RR regression was calculated as 1-cvLoss, where cvLoss is the mean normalized squared error across folds between the test data and the predictions from each fold, which is equal to the sum of squared errors divided by the total sum of squares of the target data.

Long timescale correlations and regression

To assess how pairwise correlations and prediction performance decays with time, we performed the above analyses with time shifted stimulus presentations ([Figure 5](#)). First, we ordered all eligible stimulus presentations by time and then, for each time shift, we shifted the response for each neuron in the target population by 0–5 stimulus presentations and repeated the analyses described above, i.e., we split source and target populations (20 times) and calculated pairwise correlations for pairs of neurons in MT, in SC, and for MT-SC pairs, and prediction performance within and across areas using RR regression. This procedure was performed separately for attend in and attend out trials. This revealed an expected decay in prediction performance and pairwise correlations with time. To compare these results across sessions, we calculated the percentage drop in correlations and prediction performance relative to time shift 0 (when the stimulus presentations were aligned) and compared the rate of correlation decay and rate of prediction performance decay ([Figures 5E and 5F](#)). We compared within area prediction performance decay with the pairwise correlation decay for that area, and across area prediction performance decay with correlation decay for MT-SC pairs.

Cross-condition, cross-validated regression

To assess the effect of attention on the structure of the shared subspace between interaction populations of neurons, we calculated how well the regression weight matrix for one condition (attend in, say) predicted the responses of the target population in the opposite condition (attend out). In the first analysis, we simply used the cross-validated optimum number of dimensions to obtain a weight matrix in one condition and tested it against the trials of the other condition. The results of this method are depicted in [Figures S7A–S7D](#). The accuracy of the inter-areal interaction dropped significantly but the accuracy of the intra-areal interaction was not affected. To assess whether this was a result of a linear scaling of the weight matrix across conditions due to non-stationarities or other task/stimulus independent factors, we projected the response of the source population using the weight matrix of the opposite condition before performing RR regression to obtain the prediction. This was cross-validated in the following way described in pseudo-code (for the MT → SC interaction, for the attend out trials using the attend out versus attend in models, but we followed the same process for all potential permutations of conditions and areas).

For each fold, run 1-7:

```
1. W_out = regress(MTout, train -> SCout, train)
2. SCout, testPred = predict(MTout, test, W_out)----- (A)
3. W_in = regress(MTin, train -> SCin, train)
4. MTout, train' = project(MTout, train, W_in)
5. MTout, test' = project(MTout, test, W_in)
6. W_outCross = regress(MTout, train' -> SCout, train)
7. SCout, testPredCross = predict(MTout, test', W_outCross)---(B)
attendOut_NSE_within = NSE(SCout, testPred, SCout, test)
attendOut_NSE_cross = NSE(SCout, testPredCross, SCout, test)
ratio = attendOut_NSE_cross/attendOut_NSE_within
```

The ratio thus obtained was a cross-validated measure of how well the attend out weight matrix (W_{out}) performs compared to the weight matrix ($W_{outCross}$) that is trained to predict the same activity projected through the attend in weight matrix (W_{in}) first. We ran this for 10-folds for each random split of each population (described above) and evaluated the ratio of the normalized square error of prediction using both the matrices. This ratio is a quantitative measure of how well the cross-condition weight matrix performs relative to the within-condition weight matrix and values substantially lower than 1 would indicate a drastic drop in performance and, therefore, that the linear communication subspace between the two interacting populations is qualitatively different in their structure. We found this to be true for inter-areal interactions but not within-area interactions ([Figures S7E–S7H](#)).

Factor Analysis

We used factor analysis (FA) to assess the dimensionality of neural activity within an area. FA is a static dimensionality reduction technique that does not assume the same noise variance for all recorded neurons and calculates the dimensions of greatest covariance (instead of variance). As with RR regression, the details of this analysis can be found in previous publications.^{101,129,130} We followed the same steps as previously published work to estimate the dimensionality: (1) we found the number of dimensions m_{peak} that maximized the cross-validated log-likelihood of the observed residuals; (2) we fitted a FA model with m_{peak} dimensions and chose m , using the eigenvalue decomposition, as the smallest dimensionality that captured 95% of the variance in the shared covariance matrix. These population dimensions (m) and predictive dimensions as determined from RR regression are determined by different techniques and therefore, wherever applicable, we have used these techniques to evaluate only the change of dimensionality (private or shared) between the two attention conditions instead of comparing absolute values.