Current Biology, Volume *31*

Supplemental Information

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

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

Figure S1 – related to Figure 2:

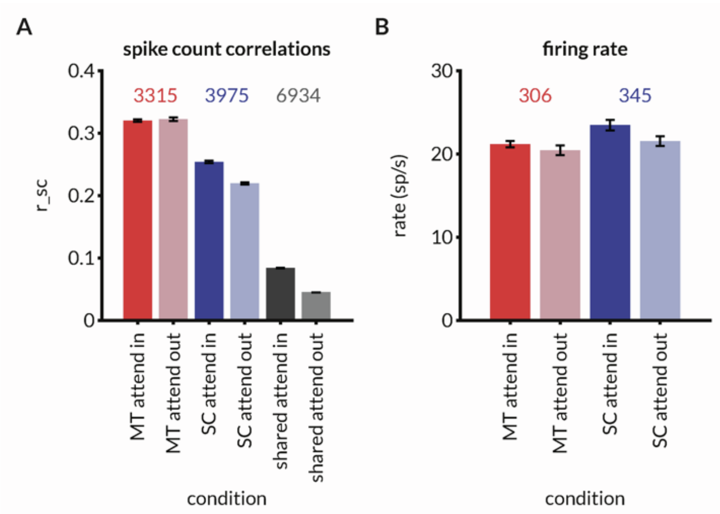

Effect of attention on aggregate noise correlations and firing rates for all neurons and pairs across all recording sessions. Error bars are standard error of the mean.

A: Spike count correlations (r_{SC}) for 3315 MT neuron pairs (red), 3975 SC neuron pairs (blue), and 6934 MT-SC pairs (gray) for attend in and attend out conditions. r_{SC} was calculated as the Pearson correlation between spike counts during all identical stimulus presentations except the first presentation after the beginning of the trial. Attention increases spike count correlations in SC pairs ($p=2.7\text{x}10^{-69}$; Wilcoxon signed rank test) and MT-SC pairs ($p=9.1\ \text{x}10^{-224}$; Wilcoxon signed rank test) and has no effect on MT pairs (p=0.8; Wilcoxon signed rank test). The disparity between these results and previously published results[1] is largely due to the selection of stimulus presentations. Here, we chose all presentations in a trial to increase statistical power in regression and factor analyses, whereas previous publications chose only the stimulus presentation before the orientation change to compare r_{SC} with behavioral outcomes.

B: Average firing rate across all presentations for 306 MT neurons (red) and 345 SC neurons (blue). Attention significantly increases firing rates of neurons in both MT ($p=8.87\text{x}10^{-14}$; Wilcoxon signed rank test) and SC ($p=5.88\text{x}10^{-42}$; Wilcoxon signed rank test).

Figure S2 – related to Figure 4:

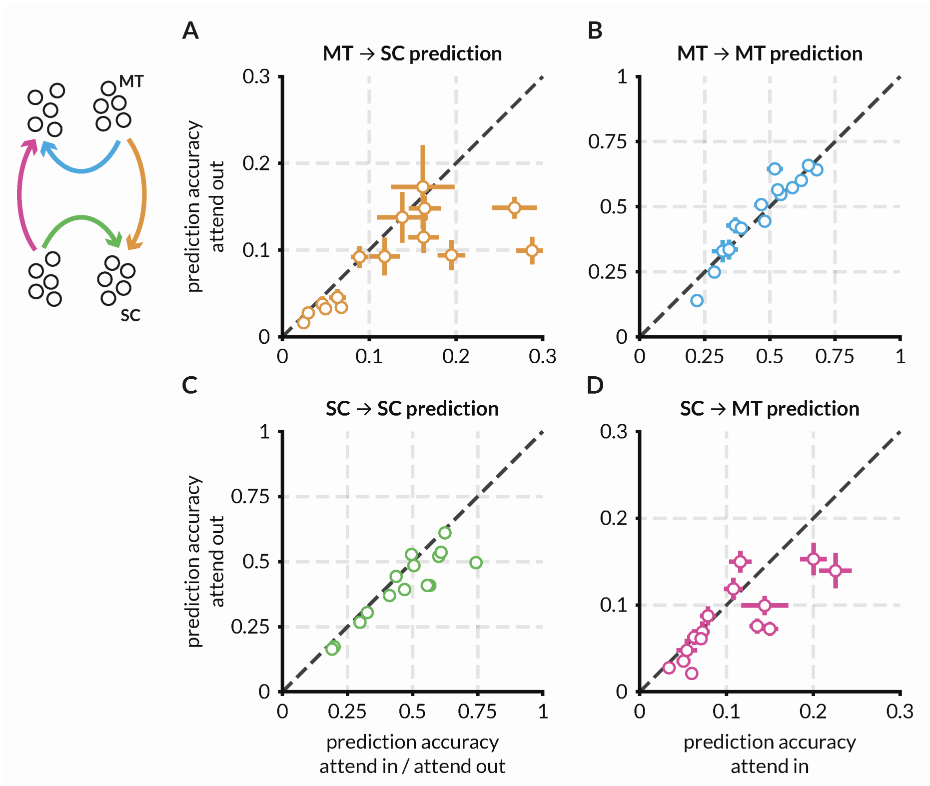

Attention improves prediction accuracy inter-areal communication. Each point represents a recording session, and the color scheme is the same as other figures and redundant with the plot labels. Error bars represent the 95% confidence intervals across all random splits of the source and target populations. Note the difference in scales in panels A/D and B/C.

a: Prediction accuracy for attend in and attend out conditions for the prediction of SC activity from MT activity. Each dot represents the average prediction accuracy and average predictive dimensions across 20 predictions of a random half of the SC population predicted by a random half of the MT population in that session. Attention significantly improves MT ➔ SC predictive performance (p=0.002; Wilcoxon's signed rank test). Average attentional modulation index (as a proxy for the size of the effect of attention on prediction accuracy) across sessions calculated as the difference of mean prediction accuracies divided by the sum was 0.164±0.04 (mean±SEM).

B: Same as (A) but for MT ➔ MT predictions. Each dot represents the average prediction accuracy and average predictive dimensions across 20 predictions of a random half of the MT population predicted by the other half of the same population in that session. Attention has no effect on prediction accuracy (p=0.89; Wilcoxon's signed rank test). Average attentional modulation of prediction accuracies was 0.005±0.019.

C: Same as (B) but for SC ➔ SC predictions. Attention has a significant effect on the prediction accuracy (p=0.002; Wilcoxon's signed rank test). Average attentional modulation of prediction accuracies was 0.068±0.02.

D: Same as (A) but for SC ➔ MT predictions. Attention increases prediction accuracy of SC ➔ MT predictions (p=0.032; Wilcoxon's signed rank test). Average attentional modulation of prediction accuracies was 0.124±0.04.

Figure S3 – related to Figure 4 and Figure 5:

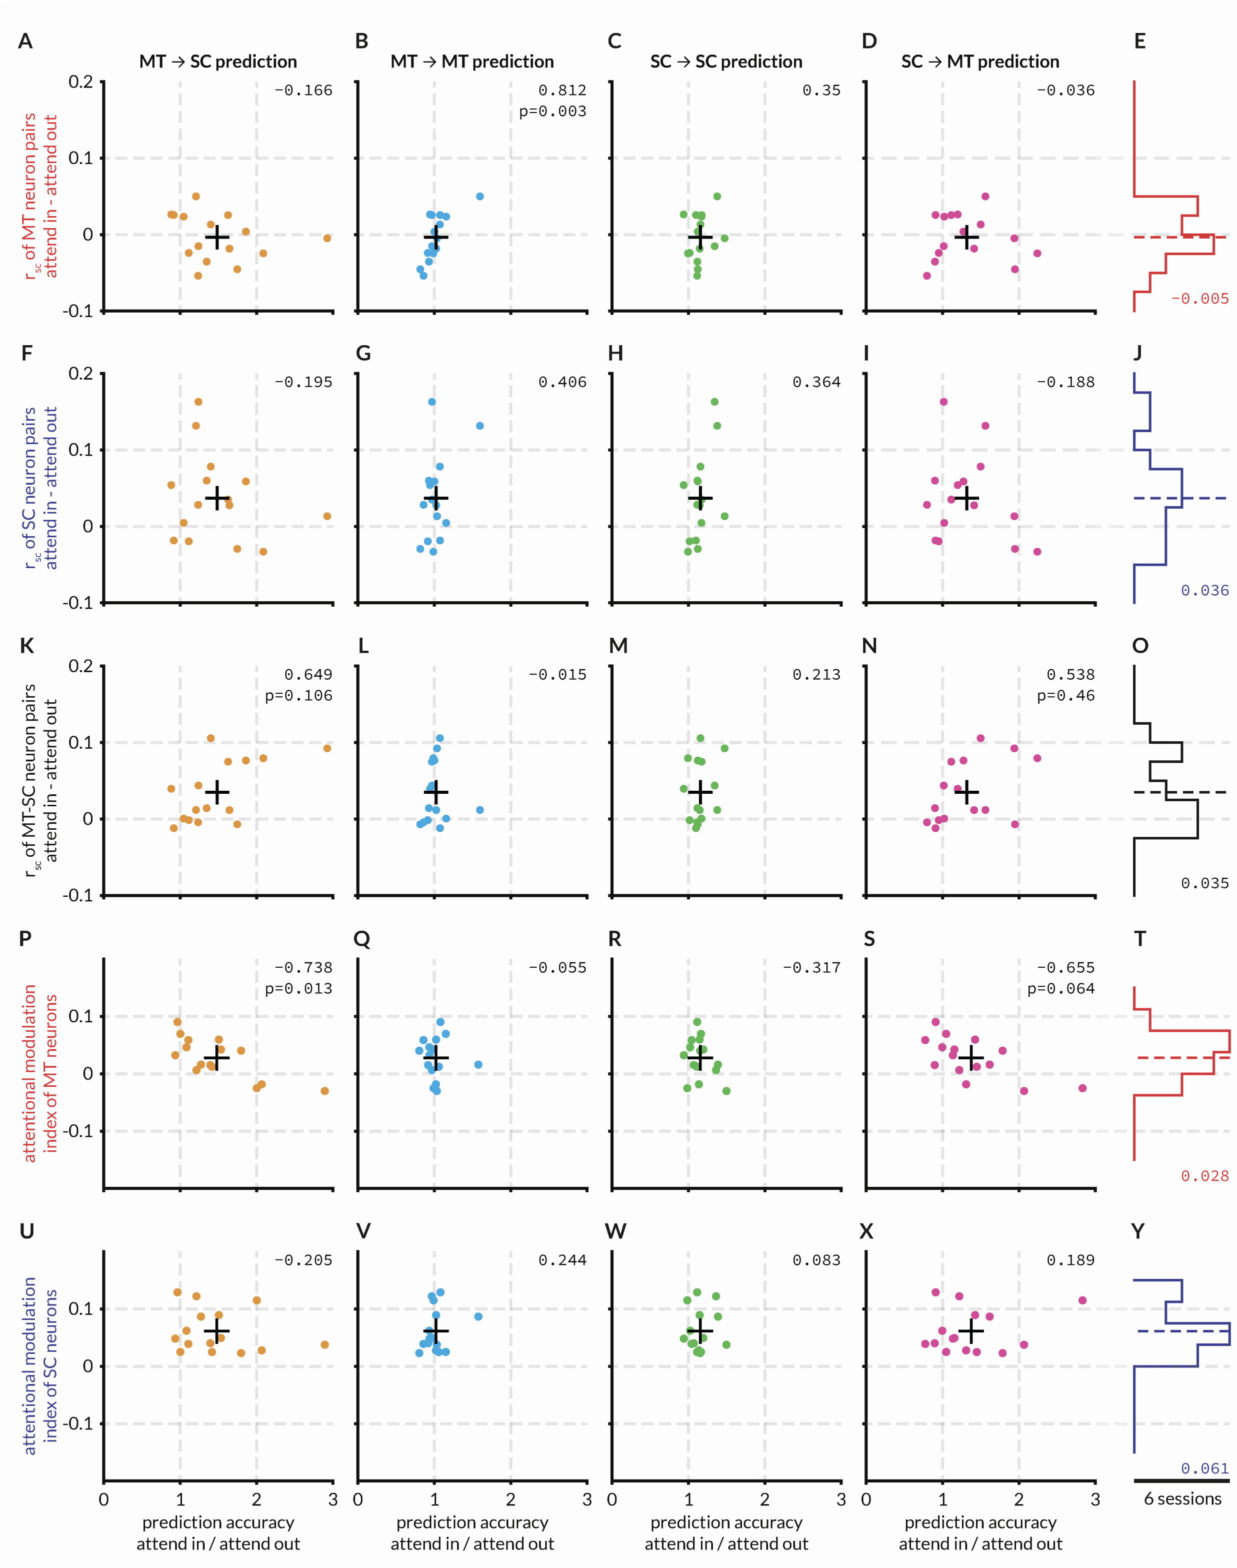

Attention-related changes in spike count correlations or attention index do not predict the improvement in communication efficacy across areas. Panels (A-O) illustrate how the differences of noise correlations of MT neuron pairs (A-E), SC neuron pairs (F-J), and MT-SC neuron pairs (K-O) between attend in and attend out conditions relate to the ratio of accuracies for within and across area response predictions. Panels (P-Y) illustrate the relationship between the average attentional modulation index for each session across neurons with the ratio of prediction accuracies between attend in and attend out conditions. Each point represents a recording session, and the color scheme is the same as other figures and redundant with the plot labels. + represents the mean of the points. The correlation coefficient for the relationship is mentioned at the top-right of each panel. For relevant relationships, the p-value is also mentioned (corrected for multiple comparisons).

A: No relationship between the effect of attention on the average accuracy of MT ➔ SC predictions for each session and the effect on the average spike count correlations for MT neuron pairs for the same session.

B: Same as (A) for MT ➔ MT predictions.

C: Same as (A) for SC ➔ SC predictions.

D: Same as (A) for SC ➔ MT predictions.

E: Histogram of the difference of spike count correlations of MT neuron pairs between the two attention conditions. Dotted line represents the mean of -0.005.

F-G: Same as A-E, but for comparing prediction accuracies with session-wise average spike count correlations for SC neuron pairs. Dotted line in the histogram in (G) represents the mean of 0.036.

K-O: Same as a-e, but for comparing prediction accuracies with session-wise average spike count correlations for MT and SC neuron pairs. Dotted line in the histogram in (O) represents the mean of 0.035. A weak relationship may be observed in (K) and (N).

P-T: Comparing the average attentional modulation for MT neurons for each session with the ratio of prediction accuracies for each regression direction. Dotted line in the histogram in (T) represents the mean MT attentional modulation across sessions of 0.028. A negative relationship between prediction performance and MT attentional modulation index may be observed in (P) and (S).

U-Y: Same as P-T, but for comparing the average attentional modulation index of SC neurons per session with the ratio of prediction performance across attention conditions. Dotted line in the histogram in (Y) represents the mean SC attentional modulation across sessions of 0.061.

Figure S4 – related to Figure 4:

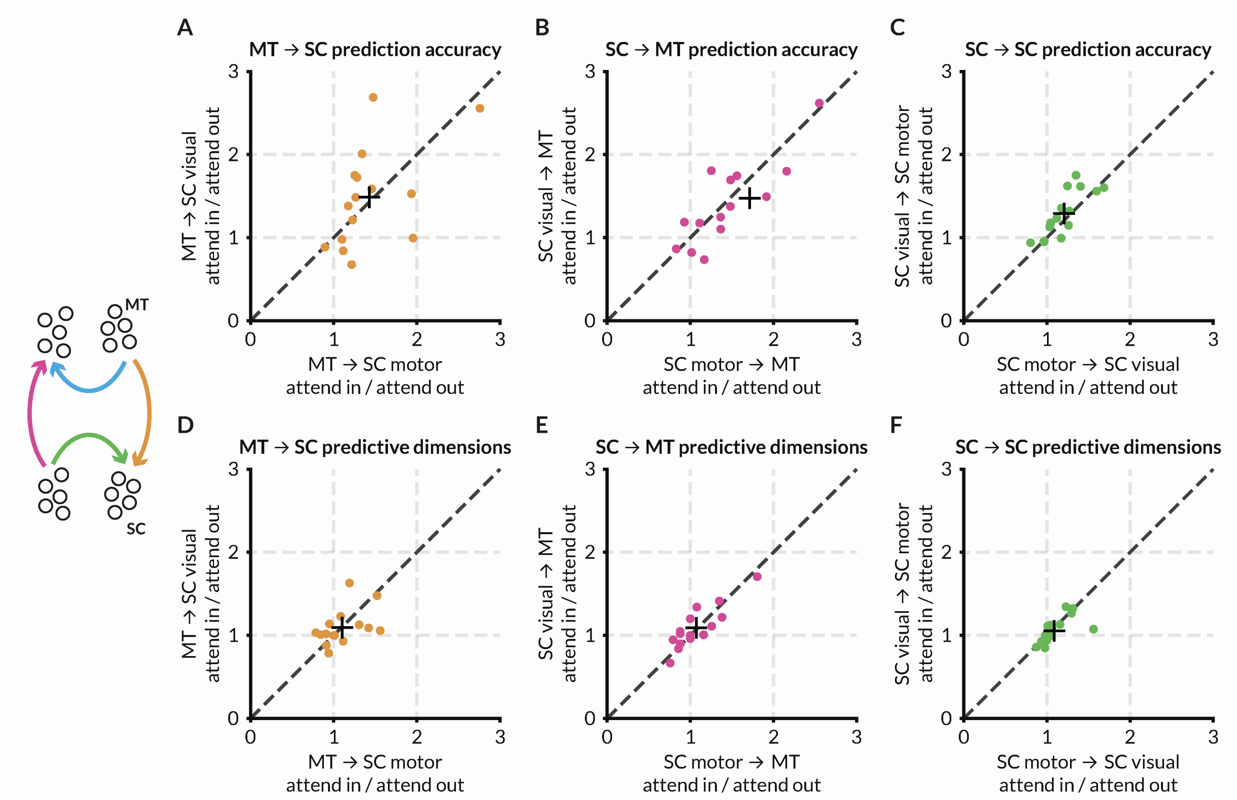

Both oculo-motor and motor neurons in SC contribute similarly to the attention-related improvement in prediction performance between MT and SC. For each session, SC neurons were ordered by an oculo-motor score (described in text and methods) and split evenly into "SC visual" and "SC motor" populations. (Oculo-motor SC neurons are labeled "SC visual" for brevity.) Each point represents a recording session, and the color scheme is the same as other figures and redundant with the plot labels. + represents the mean of the points.

A: Average accuracy of predictions of randomly split SC populations of either oculo-motor neurons or motor neurons from the same population of randomly sampled MT populations presented as a ratio of the two attention conditions. (In each iteration, 50% of randomly sampled (without replacement) MT neurons were used to predict 50% of randomly sampled SC neurons from the top half of the oculo-motor index distribution and 50% of randomly sampled SC neurons from the bottom half of the distribution. So, effectively, only 25% of the SC neurons were used for predictions in these regressions as compared to 50% in other analyses.) The prediction accuracy of both oculo-motor SC and motor SC neural activity from MT neuron activity is similarly elevated with attention. ($p = 0.0031$ for MT ➔ SC motor, $p = 0.0071$ for MT ➔ SC visual, $p = 0.52$ for the ratio of the two; one-sample t-test for the ratios)

B: Same as (A) for SC oculo-motor or motor ➔ MT predictions. As with (A), prediction accuracy is similarly enhanced with attention. ($p = 0.0309$ for SC motor ➔ MT, $p = 0.0052$ for SC visual ➔ MT, $p = 0.456$ for the ratio of the two; one-sample t-test for the ratios)

C: Same as (A) for recurrent connections between SC oculo-motor and SC motor populations. As with (A), prediction accuracy is enhanced with attention. ($p = 0.0047$ for SC motor ➔ SC visual, $p = 0.0013$ for SC visual ➔ SC motor, $p = 0.0495$ for the ratio [SC visual ➔ SC motor] / [SC motor ➔ SC visual])

D: Same as (A) but for the ratio of the average number of predictive dimensions between the two attention conditions for the MT ➔ SC oculo-motor or SC motor predictions. Attention has no effect on the dimensionality of the shared subspace between MT and SC populations. ($p > 0.05$ for all ratios; t-test)
E: Same as (B) for the ratio of the average number of predictive dimensions between the two attention conditions for the SC oculo-motor or SC motor predictions ➔ MT predictions. ($p > 0.05$ for all ratios; t-test)
F: Same as (C) for the ratio of the average number of predictive dimensions between the two attention conditions for the recurrent connections between the SC oculo-motor and SC motor populations. ($p > 0.05$ for all ratios; t-test)

Figure S5 – related to Figure 6:

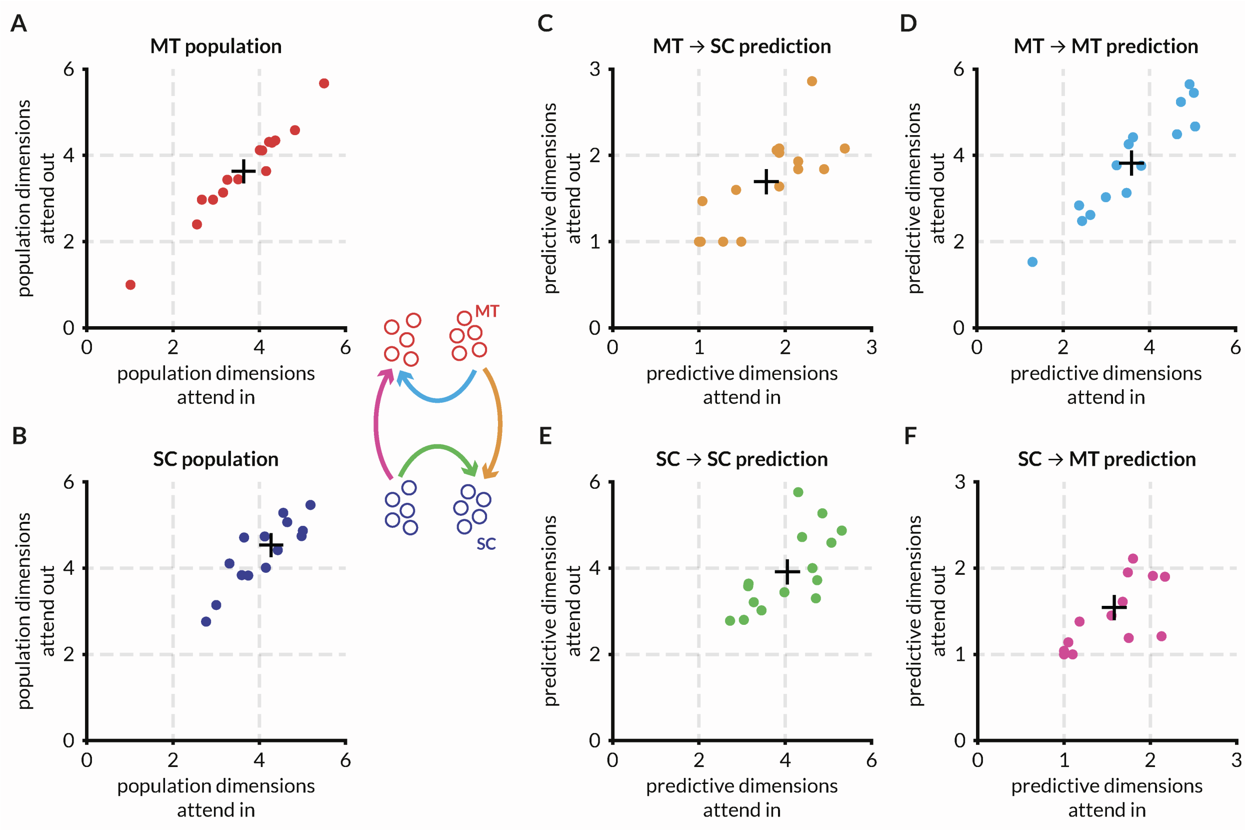

Attention does not alter the dimensionality of the response space in SC or MT, or the dimensionality of the shared communication subspace. Each point represents a recording session, and the color scheme is the same as other figures and redundant with the plot labels. + represents the mean of the points.

A: Attention does not affect the population dimensionality of the MT populations. Each point represents the average number of dimensions (factors) required to explain 95% of the variance in the MT activity for one session. On average, fluctuations in MT activity are largely restricted to ~ 3.5 dimensions.

B: Attention does not affect the population dimensionality of the SC populations. Same as (A) for the SC population. On average, fluctuations in SC activity are largely restricted to ~ 4.2 dimensions.

C-F: Attention does not affect the number of dimensions required to optimally predict target activity for any of the four predictions. Same data as Figure 4A split into four panels for clarity.

Figure S6 – related to Figure 6:

(full page width – 2 column)

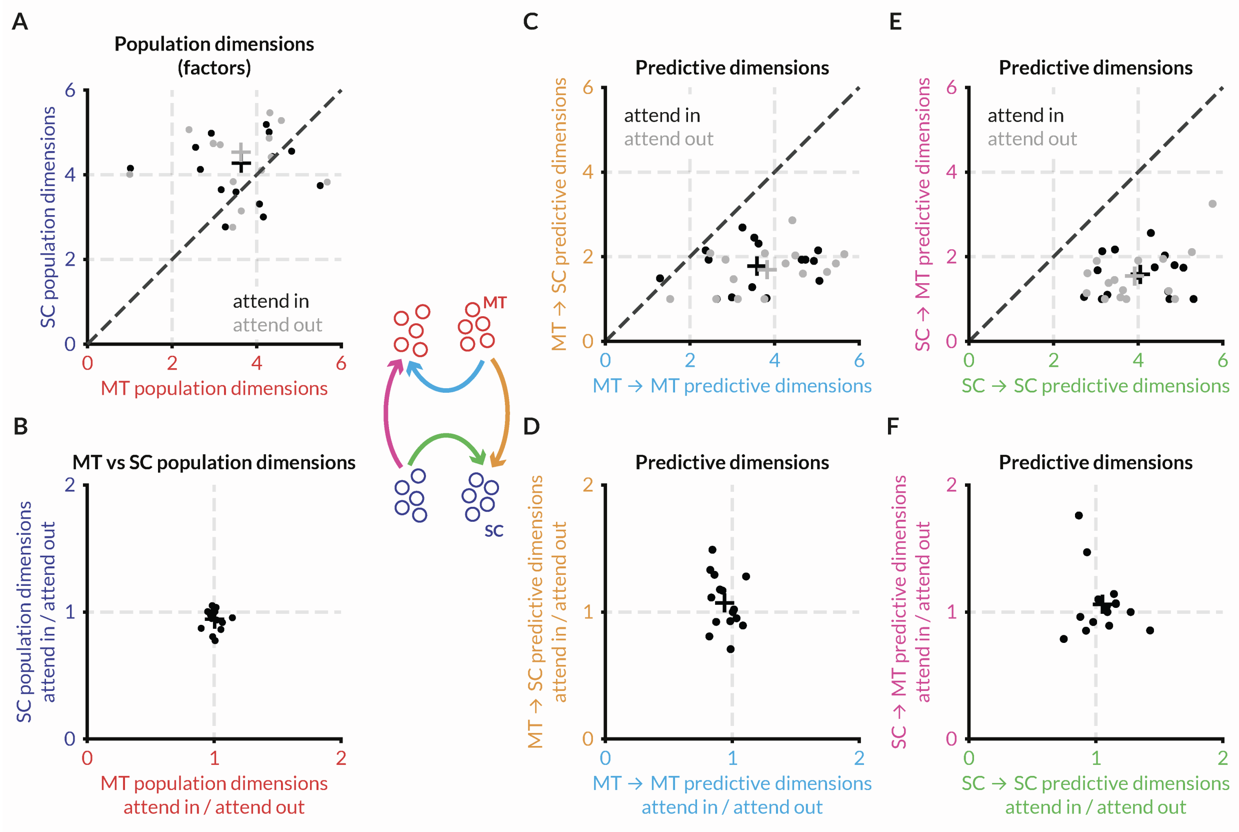

Detailed comparison of attention-related changes in MT and SC population dimensions and predictive dimensions different predictions. Each point represents a recording session, and the color scheme is the same as other figures and redundant with the plot labels. Colored + represents the mean of the corresponding points.

A: Number of population dimensions or factors from factor analysis for the MT and SC populations in each session for attend in and attend out conditions. 95% of the variance in the MT and SC population activity can be explained with approximately 3.5 and 4.3 dimensions respectively in both attention conditions.

B: Same as (A) expressed as a ratio of population dimensions in attend in and attend out conditions. Attention has no effect on the number of dimensions required to explain 95% of the variance in activity in this dataset.

C: Number of predictive dimensions that are "shared" between MT and SC (orange axis) vs the number of dimensions that are "private" in MT (blue axis) in the two attention conditions. The number of MT dimensions required to predict SC activity (~ 2) is lower than the number of MT dimensions required to predict MT activity (~ 4).

D: Same as (C) expressed as a ratio of predictive dimensions in attend in and attend out conditions.

E: Same as (C) but for the number of dimensions in SC population activity that is sufficient to explain MT activity. Number of dimensions "shared" between SC and MT (~ 2) in SC activity is lower than the number of "private" SC dimensions (~ 4).

F: Same as (E) expressed as a ratio of predictive dimensions in attend in and attend out conditions.

Figure S7 – related to Figure 4 and 6:

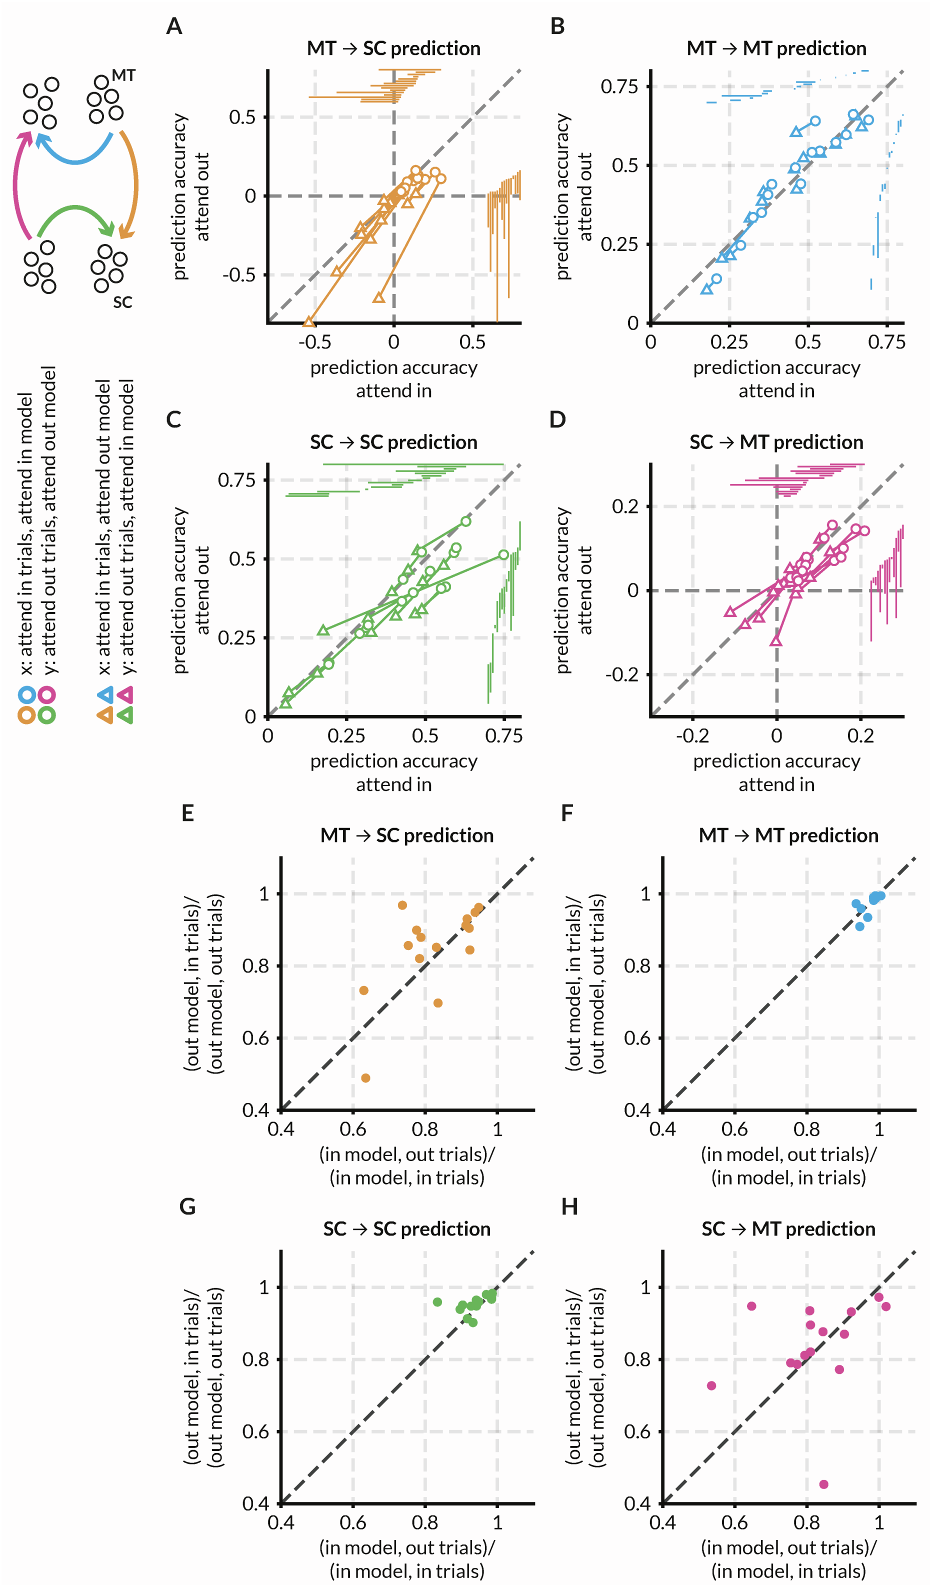

Cross-predicting activity for attend in trials using attend out model and vice versa reveals that the linear subspaces for across area communication are not identical. While the dimensionality of the communication subspace is not affected by attention, it is possible that the structure of the subspace changes while keeping its dimensionality, in turn causing the prediction accuracy to be better. To test this hypothesis, we used the weights of the linear model that corresponded to the optimum number of predictive dimensions in the attend in condition and used it to predict the target responses in the attend out condition and vice versa. We observed a marked drop in performance for cross-prediction for inter-areal communication in both directions but not intra-areal communication (A-D). To test whether this drop was due to a linear scaling of the weights across conditions and to cross-validate the cross-predictions, we projected the source activity through the weight matrix of the opposite attention condition and then fit a linear model to the target activity (see Methods for the details of the algorithm) and plotted the cross-validated cross-prediction performance normalized by the cross-validated performance of the true model. We observed a reduction in performance for the inter-areal predictions, albeit milder than earlier estimates (E-F). The intra-areal communication channels remained unaffected. While it may be possible that inter-areal communication indeed utilizes a different assortment of shared dimensions across attention conditions, we assert that these linear methods afford us a partial view of the effect of attention on the communication between areas. Each point represents the mean prediction accuracy of a recording session, and the color scheme is the same as other figures and redundant with the plot labels.

A: We plotted the average cross-prediction accuracy (triangles) for each session and each communication channel across random splits against the true prediction accuracy (circles) i.e., the cross-validated prediction accuracy of the attend in models with the attend in trials etc. The linear model trained to predict SC activity using MT responses in the attend in condition performs significantly ($p = 2.62 \times 10^{-4}$; Wilcoxon rank sum test) worse when used to predict the SC responses for trials in the opposite attend out condition; the same is true for the reverse – using the attend out model to predict the attend in responses ($p = 2.33 \times 10^{-5}$; Wilcoxon rank sum test). Circles represent mean cross-validated prediction accuracy across random splits MT and SC neurons (same points as Figure 4A). For each random split, the linear model of the opposite set of trials was used to predict the responses; the mean accuracy this out-of-set prediction across all random splits is represented by the triangles. Each circle-triangle pair is connected by a line and represents the change in prediction performance for a single session. The projections of each line on the cardinal axes are shown on the top and right of the plot, ordered by the prediction accuracy. Out-of-set prediction accuracies are always lower ($p = 2.62 \times 10^{-4}$; Wilcoxon rank sum test) and not significantly different from 0 ($p = 0.07$; t-test), which may mean that the model is unable to do better than guessing the target variance based on the mean of the target population activity (see Semedo et al., 2019 for more details[2]). Both out-of-set models are similarly affected, evident from the consistent slope of the lines. This drastic drop in performance suggests that the shared communication subspace between MT and SC is different across attention conditions.

B: Out-of-set mean accuracies for the MT ➔ MT prediction are not significantly different ($p = 0.68$ for the attend in model and $p = 0.65$ for the attend out model for attend in vs attend out trials; Wilcoxon rank sum test) suggesting not only that attention does not affect prediction performance within MT, but also that the same axes of fluctuations within the MT population activity are used for communication within MT thereby using the same private communication subspace.

C: Same as (B) but for SC ➔ SC prediction. The out-of-set prediction is not significantly different ($p = 0.046$ for the attend in model and $p = 0.097$ for the attend out model for attend in vs attend out trials; Wilcoxon rank sum test).

D: Same as (A) but for SC ➔ MT predictions. The out-of-set prediction is significantly worse for both the attend in model ($p = 0.0011$; Wilcoxon rank sum test) and the attend out model ($p = 0.0016$; Wilcoxon rank sum test).

E: To control for the case where the prediction weights across conditions may be linearly scaled and thereby produce significantly worse predictions, the following procedure was used (these steps are for comparing the MT ➔ SC attend in weights with the attend out trials, but the same procedure applies for all possible permutations of conditions and populations). The pseudo-code for this cross-validated cross-prediction method can be found in Methods. First, the MT ➔ SC prediction weights were found for a set of attend out training trials (`W_out`) and the SC activity was predicted for the test trials (`SCout,testPred`). Similarly, the prediction weights for the training set of attend in trials was found (`W_in`). Then `W_in` was used to project the attend out MT activity for both training and test trials and then used to predict SC activity in the attend out condition for the test trials (`SCout,testPredCross`). After finding predictions across all folds, the normalized square error was found and compared for the within and across condition predictions. The ratio of the across/within condition prediction for the attend in trials for each session is plotted against the ratio of the across/within condition prediction for the attend out trials. This comparison between these variables demonstrates the ability of the same communication subspace being applied to the trials in the opposite condition and therefore a ratio substantially lower than 1 would indicate that the populations communicate using different subspaces in the different conditions. The cross-prediction accuracy is significantly lower for both attend in and attend out models tested with attend out and attend in trials respectively.

F: same as (E), but for MT ➔ MT interactions. As in (B), the performance of the model from the opposite condition does not reduce prediction performance significantly.

G: same as (E), but for SC ➔ SC interactions. While the cross-prediction accuracy was not significantly different across the two attention conditions in (C), the performance of the model was lower in each session. Here, the cross-validated cross-performance shows little difference in the ratio, which provides more evidence for the hypothesis that attention does not alter the dimensionality or structure of the SC-SC communication subspace.

H: same as (E), but for SC ➔ MT interactions. As in (E), SC ➔ MT cross-prediction accuracy is significantly lower for both attend in and attend out models tested with attend out and attend in trials respectively. This difference in the structure or the constitution of the communication subspace between MT and SC between attention conditions may be evidence for attention either (a) altering the weights of interareal communication at a fast trial-to-trial timescale by unknown mechanisms, or (b) the inability of linear methods like FA and RR regression to describe potentially non-linear response spaces and the non-linear dynamics of intra- and inter-areal interactions.

Supplemental References

1. Ruff, D.A., and Cohen, M.R. (2019). Simultaneous multi-area recordings suggest that attention improves performance by reshaping stimulus representations. Nature Neuroscience *22*, 1669–1676.

2. Semedo, 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.