Published in final edited form as:

Science. 2018 January 26; 359(6374): 463–465. doi:10.1126/science.aao0284.

Learning and Attention Reveal a General Relationship Between Population Activity and Behavior

A. M. Ni, D. A. Ruff, J. J. Alberts, J. Symmonds, and M. R. Cohen*

Department of Neuroscience and Center for the Neural Basis of Cognition, University of Pittsburgh, Pittsburgh, PA 15260, USA.

Abstract

Prior studies have demonstrated that correlated variability changes with cognitive processes that improve perceptual performance. We tested whether correlated variability covaries with subjects' performance, whether performance improves quickly with attention or slowly with perceptual learning. We found a single, consistent relationship between correlated variability and behavioral performance, regardless of the time frame of correlated variability change. This correlated variability was oriented along the dimensions in population space used by the animal on a trial-by-trial basis to make decisions. That subjects' choices were predicted by specific dimensions that were aligned with the correlated variability axis clarifies longstanding paradoxes about the relationship between shared variability and behavior.

One Sentence Summary:

Correlated variability is oriented along the dimensions of neuronal population activity that best predict choices on a trial-by-trial basis.

The responses of pairs of neurons to repeated presentations of the same stimulus are typically correlated (1, 2) (quantified as noise correlations, or, spike count correlations (r_{SC})). Prior electrophysiological studies have shown that these correlations change with cognitive processes that affect perceptual performance (2–4). However, theoretical work has suggested that this correlated variability may not affect the information encoded by a neuronal population in a manner that influences a subject's decisions (5, 6).

We therefore measured the relationship between neuronal population activity and performance by studying two processes that both improve visual performance but on very different timescales: attention (7) and perceptual learning (8). By observing attention and learning in the same behavioral trials and neuronal populations, we identified the dimensions of population activity that matter most for behavior.

We recorded from neuronal populations in V4 (3, 4, 7–9) in two rhesus monkeys with chronically implanted microelectrode arrays (3; fig. 1A). The monkeys detected changes in

*Corresponding author. cohenm@pitt.edu.

The authors declare no competing financial interests. Data analyzed in this manuscript are available at <https://pitt.box.com/v/NiRuffAlbertsSymmondsCohen2017>.

the orientation of either of two Gabor stimuli, one placed within the receptive fields (RFs) of the recorded neurons and one in the opposite hemifield (fig. 1B). We measured attention effects within a single session and learning effects across sessions (fig. 1C).

Attention and perceptual learning improved performance and affected neuronal population responses in similar ways (fig. 2, S1, S2). Both processes were associated with decreases in the mean-normalized trial-to-trial variance (Fano factor) of individual units and the correlated variability between pairs of units (fig. 2, C, D, J, and K) in response to repeated presentations of the same stimulus (fig. S3, S4). These variability changes occurred only in the context of the task (variability measured during passive fixation was constant throughout training; fig. 2, F, G, M, and N).

Recent theoretical work suggests that only correlated variability along the dimensions in neuronal population space that encode task-relevant stimulus information can limit information coding (5, 6). Determining whether correlated variability lies along these dimensions is experimentally unfeasible because it would require recordings from a very large number of neurons over an even larger number of trials.

Instead, we assessed the importance of attention- and perceptual learning-related changes in correlated variability by investigating their relationship to behavior. There was a single, robust relationship between correlated variability and perceptual performance, whether changes in performance happened quickly (attention; fig. 3, A and B) or slowly (learning; fig. 3, C and D). This relationship was robust even when we removed the main effects of attention and learning (fig. 3, E and F).

We analyzed the responses of V1 neurons (7, 8) in animals performing the same attention task. Unlike in V4, correlated variability in V1 was not correlated with performance (fig. S5).

Both attention and perceptual learning improved the performance of a cross-validated, optimal linear stimulus decoder (fig. S6). However, the relationship between correlated variability in V4 and performance (fig. 3) seems at odds with theoretical work suggesting that most correlated variability should not affect the stimulus information that can be gleaned from an optimal decoder (6).

To examine the relationship between correlated variability and performance more directly, we developed a single trial measure of correlated variability. We performed principal component analysis (PCA) on population responses to the same repeated stimuli used to compute spike count correlations (fig. S3), meaning that the first PC is by definition the axis that explains more of the correlated variability than any other dimension (x-axis in fig. 4, A and B). Consistent with the recent observation that correlated variability is typically low dimensional (10–12), the variance explained by the first PC was strongly related to the magnitude of correlated variability in each session, even when we accounted for the changes caused by attention and learning (fig. 4, C and D; fig. S7) and trial-averaged firing rates (fig. S8–11). Like correlated variability (fig. 3), the proportion of variance explained by the first PC was correlated with behavioral performance (d') across all sessions (Monkey 1: $R = -0.42, p < 10^{-13}$; Monkey 2: $R = -0.62, p < 10^{-15}$).

These analyses show that projection on this first PC is a suitable proxy for pairwise spike count correlations. We used this measure to assess the importance of correlated variability to the monkey by determining whether population activity along this first PC can predict the monkey's choices on a trial-by-trial basis.

Activity along this first PC (and therefore correlated variability) had a much stronger relationship with the monkey's behavior than it would if the monkey used an optimal stimulus decoder. A linear, cross-validated choice decoder (fig. 4A) could detect differences in hit- vs. miss-trial responses to the changed stimulus as well from V4 population activity along the first PC alone as it could from our full data set (fig. 4, E and F; fig. S12). In contrast, while the performance of the stimulus decoder (fig. 4A) at detecting differences in V4 neuronal population responses to the previous stimulus (the stimulus prior to the change) vs. the changed stimulus was unsurprisingly better overall than the performance of the choice decoder (fig. 4, E and F insets), the relative influence of the first PC was weaker. The performance of the stimulus decoder was much worse when based on the first PC alone versus our full data set (fig. 4, E and F).

It is difficult to determine from extracellular recordings whether choice-predictive signals come from a bottom-up, causal relationship between sensory responses and decisions or from trial-to-trial variability from cognitive factors or post-decision signals (13). A recent study identifying the directionality of choice-predictive signals in mouse sensory cortex found that they are both bottom-up and top-down in origin (14). However, the time course of the choice- predictive activity in our data suggests that it occurs before the decision is made. We based our choice decoder on the first 70 ms of the evoked responses (after accounting for the response latency of V4 neurons). Choice-predictive activity was as strong in the first half of this time frame (60–95 ms) as in the second half (96–130 ms; paired t-test per monkey: $p > 0.05$). That the choice-predictive activity described here was present during the full decision-making period suggests that it did not reflect post-decision feedback.

Our results, combined with functional imaging in humans (8) and other multielectrode recording studies (15, 16), suggest that learning is best studied by focusing on populations of neurons. Functional imaging studies, which use measures that are related to the activity of large neural populations, find consistent learning-related changes in both V1 and V4 (8, 17), as opposed to single unit studies (8). Similarly, attention studies suggest changes in population sensitivity are largely explained by cross-neuron correlations as opposed to single neuron effects (3, 4).

The robust relationship between correlated variability and perceptual performance suggests that while attention and learning mechanisms act on different timescales (fig. S13), they share a common computation. Some characteristics of this computation are informed by recent studies showing that changes in a low rank modulator can account for the attention-related changes in rate, Fano factor, and correlated variability (11, 12). Attention and learning may decrease the strength of such a modulator by changing the balance of inhibition and excitation (10), which may improve information coding and the information that is communicated downstream (18).

Our most puzzling finding is that the attention- and learning-related changes in average noise correlation were so closely linked to performance, but would likely have a minimal effect on performance if the monkeys read out visual information optimally. Similarly, a prior study found that correlations depend on training experience, but did not find a relationship between shared variability and information coding (19). Correlated variability should only affect the performance of an optimal decoder when it is aligned with the stimulus dimension being decoded (6). Therefore, the relationship between correlated variability and performance suggests that our monkeys performed suboptimally.

We thus hypothesize that sensory information is decoded in a way that is optimal for the large number of stimuli and tasks the animals encounter in their natural environment rather than the particular set of stimuli in our task. Traditionally, optimal decoders are trained to discriminate a particular set of stimuli that vary only in one stimulus dimension. This scenario implies a two step decision process: identifying the stimulus (to optimize the decoder) and then decoding it. If animals could successfully identify the stimulus, they would perform perfectly on our change-detection task.

Instead, animals may use a more general decoder that could, for example, identify the orientation of any stimulus in any task, meaning that optimal weights would be tuned and noise correlations related to all stimulus features for which the neurons are selective. Noise correlations depend on tuning similarity for all stimulus features (6). Therefore, correlated variability is likely aligned with the dimension that is decoded by a general decoder, meaning that noise correlation decreases would improve performance. Several of the studies that suggest monkeys do behave optimally are those that used multisensory stimuli (20). Determining whether there is evidence that monkeys use decoders that are optimized for diverse stimuli and tasks will be an important avenue for future work. Our results suggest that the relationship between behavior and population activity is a powerful tool for understanding neural computation.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments:

M.R.C. is supported by US NIH grants 4R00EY020844-03, R01 EY022930, and Core Grant P30 EY008098s, a Whitehall Foundation grant, a Klingenstein-Simons Fellowship, a Sloan Research Fellowship, a McKnight Scholar Award, and a grant from the Simons Foundation. A.M.N. is supported by a fellowship from the Simons Foundation. We thank Karen McCracken for technical assistance and John H. R. Maunsell and Adam Kohn for comments on a previous version of this manuscript. A.M.N., D.A.R., and M.R.C. designed the experiments. A.M.N., D.A.R., J.J.A., and J.S. collected the data, A.M.N. performed the analyses, and A.M.N. and M.R.C. wrote the paper.

References:

1. Zohary E, Shadlen MN, Newsome WT, Nature 370, 140–143 (1994). [PubMed: 8022482]
2. Cohen MR, Kohn A, Nat Neurosci 14, 811–819 (2011). [PubMed: 21709677]
3. Cohen MR, Maunsell JHR, Nat Neurosci 12, 1594–1600 (2009). [PubMed: 19915566]
4. Mitchell JF, Sundberg KA, Reynolds JH, Neuron 63, 879–888 (2009). [PubMed: 19778515]
5. Moreno-Bote R et al., Nat Neurosci 17, 1410–1417 (2014). [PubMed: 25195105]

6. Kohn A, Coen-Cagli R, Kanitscheider I, Pouget A, Annu Rev Neurosci 39, 237–256 (2016). [PubMed: 27145916]
7. Maunsell JHR, Annu Rev Vis Sci 1, 373–391 (2015). [PubMed: 28532368]
8. Watanabe T, Sasaki Y, Annu Rev Psychol 66, 197–221 (2015). [PubMed: 25251494]
9. Ruff DA, Cohen MR, Nat Neurosci 17, 1591–1597 (2014). [PubMed: 25306550]
10. Kanashiro T, Ocker GK, Cohen MR, Doiron B, Elife 6, (2017).
11. Rabinowitz NC, Goris RL, Cohen M, Simoncelli EP, Elife 4, e08998 (2015).
12. Ecker AS, Denfield GH, Bethge M, Tolias AS, J Neurosci 36, 1775–1789 (2016). [PubMed: 26843656]
13. Nienborg H, Cumming BG, Nature 459, 89–92 (2009). [PubMed: 19270683]
14. Kwon SE, Yang H, Minamisawa G, O'Connor DH, Nat Neurosci 19, 1243–1249 (2016). [PubMed: 27437910]
15. Jeanne JM, Sharpee TO, Gentner TQ, Neuron 78, 352–363 (2013). [PubMed: 23622067]
16. Yan Y et al., Nat Neurosci 17, 1380–1387 (2014). [PubMed: 25195103]
17. Jehee JF, Ling S, Swisher JD, van Bergen RS, Tong F, J Neurosci 32, 16747–16753a (2012).
18. Ruff DA, Cohen MR, J Neurosci 36, 7523–7534 (2016). [PubMed: 27413161]
19. Gu Y et al., Neuron 71, 750–761 (2011). [PubMed: 21867889]
20. Chandrasekaran C, Curr Opin Neurobiol 43, 25–34 (2017). [PubMed: 27918886]
21. Raghunathan TE, Rosenthal R, Rubin DB, Psychol Methods 1, 178–183 (1996).
22. Haefner RM, Gerwinn S, Macke JH, Bethge M, Nat Neurosci 16, 235–242 (2013). [PubMed: 23313912]

Fig. 1. Methods and behavior.

(**A**) Orientation change-detection task with cued attention (3). (**B**) RF centers of recorded units from example session (*black circles*). Gray circles illustrate Gabor locations; red circle illustrates representative RF size. (**C**) Methodology for quantifying attention- and learning-related changes in detection sensitivity (d'). Best-fitting exponential functions plotted with S.E.M. Heat map illustrates session number. Insets: psychometric curves for two example sessions.

Fig. 2. Summary of behavioral and neuronal effects of attention and perceptual learning. All changes were significantly different than 0 except where indicated (t-tests; $p < 10^{-3}$). Conventions as in 1C. (**A, H**) Sensitivity (d') increased with both attention and learning. (**B, E, I, L**) Evoked response (firing - baseline rate) increased with attention but did not change consistently with learning or passive fixation ($p > 0.05$). (**C, D, J, K**) Fano factor and correlated variability decreased with attention and learning, (**F, G, M, N**) but not during passive fixation ($p > 0.05$).

Fig. 3. The relationship between correlated variability and performance is the same for attention and perceptual learning.

(A-D) Mean r_{SC} and d' were significantly correlated across sessions ($p < 10^{-3}$). (A, B) Relationship between r_{SC} and d' was indistinguishable between attention conditions (Fisher's z Pearson-Filon tests; $p > 0.05$). (C, D) Relationship between r_{SC} and d' was indistinguishable for the first vs. second half of learning ($p > 0.05$). (E, F) Relationships persisted after removing attention and learning effects (residuals of exponential fits in fig. 2; $p < 10^{-3}$; ANOVAs: $p > 0.05$).

Fig. 4. Correlated variability is more closely aligned with choices than would be expected from an optimal stimulus decoder.

(A) Schematic showing how we obtained our single trial measure of correlated variability. We performed PCA on responses to the stimulus before the change (*gray*) and to the changed stimulus (*black*). The stimulus decoder detects differences between the responses to the previous and changed stimuli. The choice decoder detects differences between responses to misses and hits. (B) Example data set showing that the animal's choices are more aligned with the first PC (x-axis) than the difference between the previous and changed stimuli, which depends on both the first and second PC. (C, D) Mean r_{SC} is highly correlated with proportion of variance explained by the first PC ($p < 10^{-11}$; residuals: $p < 10^{-9}$). (E, F) Performance of the choice and stimulus decoders normalized to their maximum performance (with the full data set), with S.E.M. Inset: raw decoder performance.